arXiv:1610.01970v1 [math.OC] 6 Oct 2016

Adaptive Sequential Stochastic Optimization

Craig Wilson*
Google
wilson60@illinois.edu

Venugopal Veeravalli
University of Illinois at Urbana-Champaign
vvv@illinois.edu

Angelia Nedić
Arizona State University
Angelia.Nedich@asu.edu

Abstract

A framework is introduced for sequentially solving convex minimization problems where the objective functions change slowly, in the sense that the distance between successive minimizers is bounded. The minimization problems are solved by sequentially applying a selected optimization algorithm, such as stochastic gradient descent (SGD), based on drawing a number of samples in order to carry out a desired number of iterations. Two tracking criteria are introduced to evaluate approximate minimizer quality: one based on being accurate with respect to the mean trajectory, and the other based on being accurate in high probability (IHP). Knowledge of a bound on the minimizers' change, combined with properties of the chosen optimization algorithm, is used to select the number of samples needed to meet the desired tracking criterion. A technique to estimate the change in minimizers is provided along with analysis to show that eventually the estimate upper bounds the change in minimizers. This estimate of the change in minimizers provides sample size selection rules that guarantee that the tracking criterion is met. Simulations are used to confirm that the estimation approach provides the desired tracking accuracy in practice.

1 Introduction

Problems involving optimizing a sequence of functions that slowly vary over time naturally arise in many different contexts including channel estimation, parameter tracking, and sequential learning. To describe and analyze such problems, we consider solving a sequence of stochastic convex optimization problems

$$\min_{x\in\mathscr{X}} \left\{ f_n(\boldsymbol{x}) \triangleq \mathbb{E}_{\boldsymbol{z}_n}\left[\ell_n(\boldsymbol{x},\boldsymbol{z}_n)\right] \right\} \tag{1}$$

with ℓ_n being an appropriate loss function, $\boldsymbol{z}_n$ representing the randomness in the loss at time n, and $\boldsymbol{x}\in\mathscr{X}\subset\mathbb{R}^d$ being a nonempty, closed and convex set. We will assume that problem eq. (1) has a unique solution, denoted by $\boldsymbol{x}_n^*$, at every instance n, i.e.,

$$\boldsymbol{x}_n^* = \arg\min_{\boldsymbol{x}\in\mathscr{X}} f_n(\boldsymbol{x}) \qquad \text{for all } n\geq 1.$$

To capture the idea that the sequence of functions in eq. (1) is changing slowly, we assume that there is a bound $\rho>0$ on the optimal solutions of the form:

$$\|\boldsymbol{x}_{n+1}^* - \boldsymbol{x}_n^*\|_{L_2} \leq \rho \tag{2}$$

where $\|X\|_{L_q} = \left(\mathbb{E}[X^q]\right)^{1/q}$. Rather than using a Markov chain model or other Bayesian models for the changes in $\{\boldsymbol{x}_n^*\}_{n=1}^{\infty}$, we only use the bound eq. (2) in our analysis. The formulation in eq. (2) involves an expectation to allow

*This work was supported by the NSF under award CCF 1111342 and NSF DMS 1312907 through the University of Illinois at Urbana-Champaign. Part of this work was presented at CDC 2014 [1] and ICASSP2016 [2].

for the possibility that the functions $f_n(\boldsymbol{x})$ themselves evolve stochastically. In this case, the minimizer $\boldsymbol{x}_n^*$ is itself a random variable.

Given a sequence of slowly varying functions $\{f_n(\boldsymbol{x})\}_{n=1}^{\infty}$, we want to efficiently, sequentially minimize each of the functions to within a desired accuracy. We look at solving this problem by applying an optimization algorithm such as stochastic gradient descent (SGD) that uses K_n stochastic gradient steps. To be precise, suppose that we have a function $\boldsymbol{g}_n(\boldsymbol{x},\boldsymbol{z})$ such that, given $\boldsymbol{z}_n \sim p_n$ from an auxiliary distribution p_n, it holds that

$$\mathbb{E}_{\boldsymbol{z}_n\sim p_n}[\boldsymbol{g}_n(\boldsymbol{x},\boldsymbol{z}_n) \mid \boldsymbol{x}] = \nabla f_n(\boldsymbol{x}) \qquad \text{for all } \boldsymbol{x}. \tag{3}$$

For example, if the functions $\{f_n(\boldsymbol{x})\}_{n=1}^{\infty}$ are of the form

$$f_n(\boldsymbol{x}) = \mathbb{E}_{\boldsymbol{z}_n\sim p_n}[\ell(\boldsymbol{x},\boldsymbol{z}_n)] \tag{4}$$

where $\ell(\boldsymbol{x},\boldsymbol{z})$ can correspond to a loss in machine learning, then we can set

$$\boldsymbol{g}_n(\boldsymbol{x},\boldsymbol{z}_n) = \nabla_{\boldsymbol{x}}\ell(\boldsymbol{x},\boldsymbol{z}_n)$$

under suitable technical conditions. We want to understand the trade-off between the solution accuracy and the complexity, represented by the number of stochastic gradients K_n. In effect, we want to understand how many samples $\{\boldsymbol{z}_n(k)\}_{k=1}^{K_n}$ are necessary to achieve a desired level of accuracy.

We introduce two different types of tracking criteria to characterize approximate minimizers of eq. (1), denoted $\boldsymbol{x}_n$ for each n. First, we define a *mean tracking criterion*

$$\mathbb{E}[f_n(\boldsymbol{x}_n)] - f_n(\boldsymbol{x}_n^*) \leq \varepsilon \tag{5}$$

and second, we define an *in high probability (IHP) tracking criterion*

$$\mathbb{P}\{f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t\} \leq r \tag{6}$$

with the expectation and probability taken over the samples $\{\boldsymbol{z}_n(k)\}_{k=1}^{K_n}$. We mainly focus on the mean criterion criterion paper with the IHP criterion analyzed in Section 11.

The remainder of this paper is organized as follows. In Section 2, we introduce our problem. In Section 3, we study the problem of selecting the number of samples K_n to achieve either the mean criterion in eq. (5). We find a relationship between K_n and ε for the mean tracking criterion with the change in the minimizers eq. (2) known. This relationships allows us to select K_n in order to satisfy the mean criterion. In Section 4, we introduce an estimate for the change in the minimizers, ρ, from eq. (2). We provide theoretical guarantees that the introduced estimate eventually upper bounds the change in the minimizers. In Section 5, we combine the ρ estimate of Section 4 with the analysis of the case with ρ known in Section 3 to provide rules to select K_n in order to meet the desired tracking criterion. We provide guarantees that for n large enough, we meet our desired tracking criterion almost surely. Finally, we carry out simulation experiments to test our ρ estimation and K_n selection rules.

1.1 Related Work

There has been some work on similar problems, but general optimization theory tools to deal with time-varying optimization problems under eq. (2) have yet to be developed. One relevant approach is online optimization in which a sequence of functions arrive, and in general no knowledge is available about the incoming functions other than that all the functions come from a specified class of functions, i.e., linear or convex functions with uniformly bounded gradients. Online optimization models do not include the notion of a desired tracking accuracy at each time instant such as eq. (5) and eq. (6). Instead, only bounds on the worst case performance of the best estimators are investigated through regret formulations [3–12].

For the problem of online optimization, the idea of controlling the variation of the sequence of functions has been studied in [13] and [14]. In [14], regret is minimized subject to a bound, say G_b, on the total variation of the gradients over a time interval T of interest, i.e.,

$$\sum_{n=2}^{T} \max_{\boldsymbol{x}\in\mathscr{X}} \|\nabla f_n(\boldsymbol{x}) - \nabla f_{n-1}(\boldsymbol{x})\|_2^2 \leq G_b. \tag{7}$$

If all the functions $\{f_n(x)\}$ are strongly convex with the same parameter m, then by the optimality conditions (see Theorem 2F.10 in [15]) relation eq. (7) implies that

$$\sum_{n=2}^{T} \|\boldsymbol{x}_n^* - \boldsymbol{x}_{n-1}^*\|_2^2 \leq \tilde{G}_b$$

with $\tilde{G}_b$ a function of G_b. Therefore, the work in [14] can be seen as studying the regret while controlling the total variation in the optimal solutions over T time instants. In contrast, we control the variation of the optimal solutions at each time instant with eq. (2) and then seek to maintain a tracking criterion such as eq. (5) and eq. (6) at each time instant.

Additionally, there is other work that has some of the ingredients of our proposed problem formulation. In [16], a sequence of quadratic functions is considered and treated within the domain of estimation theory; however, the authors only examine the Least Mean Squares (LMS) algorithm (corresponding to $K_n = 1$ for all n). The work in [17, 18] considers a sequence $\{f_n\}$ of convex objective functions converging to some limit function f, where all the functions f_n have the same set of minima. However, aside from considering time-varying objective functions, these works have nothing else in common with the work described here. There has also been work in [19] considering the limit as the rate of change of the functions goes to zero and for the least means squares (LMS) algorithm in [20]. The results in [19] and [20] both require a Bayesian model for the changes in the function sequence, which we do not require.

If we have a quadratic loss centered at $\boldsymbol{x}_n^*$ and a linear state space evolution for the optimal solution $\boldsymbol{x}_n^*$, then we could apply the Kalman filter [21]. If the function we seek to optimize is non-linear, another approach we can consider under a Bayesian framework is particle filtering [22]. For particle filtering, it is harder to provide exact guarantees on performance similar to those given in eq. (5) and eq. (6).

To conclude, there are no existing approaches within optimization theory or estimation theory that allow us to solve a sequence of time-varying problems, subject to abiding to a pre-specified tracking error criterion such as eq. (5) or eq. (6) under only assumption eq. (2). In this work, we fill in this gap and provide methods to solve such problems.

2 Problem Formulation

2.1 Assumptions

We make several assumptions to proceed. First, let $\mathscr{X}$ be closed and convex with $\text{diam}(\mathscr{X}) < +\infty$. Define the σ-algebra

$$\mathscr{F}_i \triangleq \sigma\left(\bigcup_{j=1}^{i}\bigcup_{k=1}^{K_j} \boldsymbol{z}_j(k)\right) \tag{8}$$

where by convention $\mathscr{F}_0$ is the trivial σ-algebra. We suppose that the following conditions hold:

A.1 For each n, $f_n(\boldsymbol{x})$ is twice continuously differentiable with respect to $\boldsymbol{x}$.

A.2 For each n, $f_n(\boldsymbol{x})$ is strongly convex with a parameter $m > 0$, i.e.,

$$f_n(\tilde{\boldsymbol{x}}) \geq f_n(\boldsymbol{x}) + \langle \nabla_{\boldsymbol{x}} f_n(\boldsymbol{x}), \tilde{\boldsymbol{x}} - \boldsymbol{x}\rangle + \frac{1}{2} m \|\tilde{\boldsymbol{x}} - \boldsymbol{x}\|^2. \tag{9}$$

A.3 For each n, we can draw stochastic gradients $\boldsymbol{g}_n(\boldsymbol{x}, \boldsymbol{z}_n)$ such that eq. (3) holds.

A.4 Given an optimization algorithm that generates an approximate minimizer

$$\boldsymbol{x}_n \triangleq \mathscr{A}(\boldsymbol{x}_{n-1}, \{\boldsymbol{g}_n(\boldsymbol{x}, \boldsymbol{z}(k))\}_{k=1}^{K_n})$$

using K_n stochastic gradients $\{\boldsymbol{g}_n(\boldsymbol{x}, \boldsymbol{z}_n(k))\}_{k=1}^{K_n}$, there exists a function $b(d_0, K_n)$ such that the following conditions hold:

1. With K_n and d_0 both $\mathscr{F}_{n-1}$-measurable, it holds that

$$\|\boldsymbol{x}_{n-1} - \boldsymbol{x}_n^*\|^2 \leq d_0^2 \;\Rightarrow\; \mathbb{E}[f_n(\boldsymbol{x}_n) \mid \mathscr{F}_{n-1}] - f_n(\boldsymbol{x}_n^*) \leq b(d_0, K_n).$$

2. With K_n a constant, it holds that $\mathbb{E}[b(d_0,K_n)] = b\left(\left(\mathbb{E}[d_0^2]\right)^{1/2}, K_n\right)$.

3. The bound $b(d_0,K_n)$ is non-decreasing in d_0 and non-increasing in K_n.

A.5 There exist constants $A, B \geq 0$ such that

$$\mathbb{E}\left[\|\boldsymbol{g}_n(\boldsymbol{x},\boldsymbol{z}_n)\|_2^2 \mid \mathscr{F}_{n-1}\right] \leq A + B\|\boldsymbol{x} - \boldsymbol{x}_n^*\|_2^2 \tag{10}$$

A.6 Initial approximate minimizers $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ satisfy

$$f_i(\boldsymbol{x}_i) - f_i(\boldsymbol{x}_i^*) \leq \varepsilon_i \quad i = 1,2$$

with ε_1 and ε_2 known.

For assumption A.4, we look at SGD

$$\begin{aligned} \boldsymbol{x}_n(k) &= \Pi_{\mathscr{X}}[\boldsymbol{x}_n(k-1) - \mu_n(k)\boldsymbol{g}_n(\boldsymbol{x},\boldsymbol{z}_n(k))] \\ \boldsymbol{x}_n(0) &\triangleq \boldsymbol{x}_{n-1} \end{aligned} \tag{11}$$

with $k = 1, K_n$. Given the iterates, we choose $\boldsymbol{x}_n$ as a convex combination of the iterates $\{\boldsymbol{x}_n(k)\}_{k=0}^{K_n}$ generated by SGD

$$\boldsymbol{x}_n = \sum_{k=0}^{K_n} \lambda_n(k)\boldsymbol{x}_n(k).$$

One simple choice is setting $\boldsymbol{x}_n = \boldsymbol{x}_n(K_n)$, which corresponds to setting $\lambda_n(K_n) = 1$ and $\lambda_n(0) = \cdots = \lambda_n(K_n - 1) = 0$. Section 10 discusses several applicable bounds $b(d_0,K)$ for SGD and choices of convex combinations $\{\lambda_n(k)\}$. In practice, we may not know the parameters such as the strong convexity parameter m from assumption A.2 and the gradient parameters A and B from assumption A.5. Section 13 introduces several techniques to estimate these parameters using the stochastic gradients in assumption A.3.

In our assumptions, we condition on the σ-algebra $\mathscr{F}_{n-1}$, since this captures all of the information available at the beginning of time n. In later sections, we will select K_n as a function of the stochastic gradients $\{\boldsymbol{g}_i(\boldsymbol{x},\boldsymbol{z}_i(k))\}_{k=1}^{K_i}$ for $i = 1,\ldots,n-1$. This implies that K_n is $\mathscr{F}_{n-1}$ measurable. In this case, where K_n is itself a random variable, assumption A.4 is crucial to our analysis.

Finally, for assumption A.6, we generally must select K_1 and K_2 blindly in the sense that we have no information about ρ defined in eq. (2). We can only make a choice such as

$$K_i = \min\left\{K \geq 1 \;\middle|\; b(\operatorname{diam}(\mathscr{X}),K) \leq \varepsilon\right\} \quad i = 1,2$$

or fixed initial choices for K_1 and K_2. Regardless of our choice of K_1 and K_2, we can set $\varepsilon_i \triangleq b(\operatorname{diam}(\mathscr{X}),K)$ for $i = 1,2$. In order to have $\varepsilon_i \leq \varepsilon$ for $i = 1,2$, we may need to draw significantly more samples up front to find points $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ due to using $\operatorname{diam}(\mathscr{X})$.

2.2 Constructing a Bound On the Change in Minimizers

We look at the justification behind our choice of controlling the change in functions through the minimizers $\boldsymbol{x}_n^*$ by showing that several other reasonable ways to control how the functions change can be reduced to a bound on the change in minimizers. In Section 3, we show that bounds on the change in the minimizer can be used to select the number of stochastic gradients K_n.

2.2.1 Change in f

Suppose that we instead bound the change in the optimal function values, in the following manner:

$$f_n(\boldsymbol{x}_{n-1}^*) - f_n(\boldsymbol{x}_n^*) \leq \tilde{\rho}.$$

This bounds the loss incurred as a result of using the minimizer of a previous function f_{n-1} and the next function f_n. By the strong convexity assumption A.2, it holds that

$$\|\boldsymbol{x}_n^* - \boldsymbol{x}_{n-1}^*\|_2 \leq \sqrt{\frac{2}{m}\left(f_n(\boldsymbol{x}_{n-1}^*) - f_n(\boldsymbol{x}_n^*)\right)} \leq \sqrt{\frac{2}{m}\tilde{\rho}}.$$

Therefore, a bound on the optimal function values can be translated into a bound on the change in the minimizers.

2.2.2 Change in Distribution

As mentioned before, for learning problems in eq. (4), we can generally write our functions as an expectation of a loss function $\ell(\boldsymbol{x}, \boldsymbol{z})$, i.e., $f_n(\boldsymbol{x}) = \mathbb{E}_{\boldsymbol{z}_n \sim p_n}[\ell(\boldsymbol{x}, \boldsymbol{z}_n)]$. Therefore, the source of change in this problem is the distributions p_n. We can control the change by making an assumption on how p_n changes through an appropriate probability metric or pseudo-metric. Given a class of functions $\mathscr{F}$ mapping from $\mathscr{Z} \to \mathbb{R}$, an integral probability metric [23] between two distributions p and q on $\mathscr{Z}$ is defined as

$$\gamma_{\mathscr{F}}(p, q) \triangleq \sup_{h \in \mathscr{F}} |\mathbb{E}_{\boldsymbol{z} \sim p}[h(\boldsymbol{z})] - \mathbb{E}_{\tilde{\boldsymbol{z}} \sim p}[h(\tilde{\boldsymbol{z}})]|.$$

The following lemma shows that under an inclusion condition on the loss function $\ell(\boldsymbol{x}, \boldsymbol{z})$, the integral probability metric bounds can lead to bounds on the change in minimizers.

Lemma 1. *If the class $\{\ell(\boldsymbol{x}, \cdot) \mid \boldsymbol{x} \in \mathscr{X}\} \subset \mathscr{F}$ of loss functions is such that $\gamma_{\mathscr{F}}(p_n, p_{n-1}) \leq \tilde{\rho}$ for all $n \geq 1$, then it holds that*

$$\|\boldsymbol{x}_n^* - \boldsymbol{x}_{n-1}^*\|_2 \leq \sqrt{\frac{2}{m}\tilde{\rho}} \qquad \text{for all } n \geq 1.$$

Proof. Applying the strong convexity assumption A.2 to $f_n(\boldsymbol{x})$ and $f_{n-1}(\boldsymbol{x})$, for the solutions $\boldsymbol{x}_n^*$ and $\boldsymbol{x}_{n-1}^*$, we obtain

$$\begin{aligned} f_n(\boldsymbol{x}_{n-1}^*) &\geq f_n(\boldsymbol{x}_n^*) + \frac{1}{2}m\|\boldsymbol{x}_n^* - \boldsymbol{x}_{n-1}^*\|_2^2 \\ f_{n-1}(\boldsymbol{x}_n^*) &\geq f_{n-1}(\boldsymbol{x}_{n-1}^*) + \frac{1}{2}m\|\boldsymbol{x}_n^* - \boldsymbol{x}_{n-1}^*\|_2^2. \end{aligned}$$

By adding these two inequalities and rearranging, it holds that

$$\begin{aligned} m\|\boldsymbol{x}_n^* - \boldsymbol{x}_{n-1}^*\|_2^2 &\leq \left(f_n(\boldsymbol{x}_{n-1}^*) - f_n(\boldsymbol{x}_n^*)\right) + \left(f_{n-1}(\boldsymbol{x}_n^*) - f_{n-1}(\boldsymbol{x}_{n-1}^*)\right) \\ &= \left(f_n(\boldsymbol{x}_{n-1}^*) - f_{n-1}(\boldsymbol{x}_{n-1}^*)\right) + \left(f_{n-1}(\boldsymbol{x}_n^*) - f_n(\boldsymbol{x}_n^*)\right). \end{aligned}$$

Now, examine the term $f_n(\boldsymbol{x}_{n-1}^*) - f_{n-1}(\boldsymbol{x}_{n-1}^*)$. By relation eq. (4), we have

$$\begin{aligned} f_n(\boldsymbol{x}_{n-1}^*) - f_{n-1}(\boldsymbol{x}_{n-1}^*) &\leq |\mathbb{E}_{\boldsymbol{z}_n \sim p_n}\left[\ell(\boldsymbol{x}_{n-1}^*, \boldsymbol{z}_n)\right] - \mathbb{E}_{\boldsymbol{z}_{n-1} \sim p_{n-1}}\left[\ell(\boldsymbol{x}_{n-1}^*, \boldsymbol{z}_{n-1})\right]| \\ &\leq \sup_{f \in \mathscr{F}} |\mathbb{E}_{\boldsymbol{z}_n \sim p_n}\left[\ell(\boldsymbol{x}_{n-1}^*, \boldsymbol{z}_n)\right] - \mathbb{E}_{\boldsymbol{z}_{n-1} \sim p_{n-1}}\left[\ell(\boldsymbol{x}_{n-1}^*, \boldsymbol{z}_{n-1})\right]| \\ &= \gamma_{\mathscr{F}}(p_n, p_{n-1}). \end{aligned}$$

Similarly, we can see that the same estimate holds for the term $f_{n-1}(\boldsymbol{x}_n^*) - f_n(\boldsymbol{x}_n^*)$. Therefore, it holds that

$$\|\boldsymbol{x}_n^* - \boldsymbol{x}_{n-1}^*\|_2 \leq \sqrt{\frac{2}{m}\gamma_{\mathscr{F}}(p_n, p_{n-1})} \leq \sqrt{\frac{2}{m}\tilde{\rho}}.$$

□

Thus, we see that we can translate a bound on the change in distributions through an integral probability metric to a bound on the change in minimizers.

2.2.3 Parameterized Functions

Finally, we examine the case in which the functions $\{f_n(\boldsymbol{x})\}$ come from a parameterized class of functions $f(\boldsymbol{x},\theta)$, i.e.,

$$f_n(\boldsymbol{x}) = f(\boldsymbol{x},\boldsymbol{\theta}_n).$$

Furthermore, we assume that the parameters themselves change slowly

$$\|\boldsymbol{\theta}_n - \boldsymbol{\theta}_{n-1}\|_2 \leq \delta.$$

With appropriate assumptions on the function $f(\boldsymbol{x},\boldsymbol{\theta})$, we can apply the implicit function theorem [24] to yield a bound of the form

$$\|\boldsymbol{x}^*_{n+1} - \boldsymbol{x}^*_n\|_{L_2} \leq G\|\boldsymbol{\theta}_{n+1} - \boldsymbol{\theta}_n\|_2 \leq G\delta$$

for an appropriately chosen G.

3 Tracking Analysis with Change in Minimizers Known

In this section we combine the bound $b(d_0,K)$ in assumption A.4 with our model for the changes in functions in eq. (2) to choose the number of stochastic gradients K needed to achieve desired mean criterion ε in eq. (5). The IHP criterion in eq. (6) is analyzed in Section 11. In this section, we assume that ρ is known. In Section 5, we will consider the case when ρ is unknown.

3.1 Mean Criterion Analysis

We show how to choose K to achieve a target mean criterion ε for all n. The idea behind the analysis is to proceed by induction using assumption A.6 as a base case. Suppose that

$$\mathbb{E}[f_{n-1}(\boldsymbol{x}_{n-1})] - f_{n-1}(\boldsymbol{x}^*_{n-1}) \leq \varepsilon.$$

Denote the distance from the initial point $\boldsymbol{x}_{n-1}$ to the minimizer $\boldsymbol{x}^*_n$ by $d_n(0)$, i.e.,

$$d_n^2(0) = \|\boldsymbol{x}_{n-1} - \boldsymbol{x}^*_n\|_2^2. \tag{12}$$

To bound $\mathbb{E}[d_n^2(0)]$ we first use the triangle inequality and eq. (2) to get

$$\begin{aligned}\sqrt{\mathbb{E}[d_n^2(0)]} &\leq \|\boldsymbol{x}_{n-1} - \boldsymbol{x}^*_{n-1}\|_{L_2} + \|\boldsymbol{x}^*_{n-1} - \boldsymbol{x}^*_n\|_{L_2} \\ &\leq \|\boldsymbol{x}_{n-1} - \boldsymbol{x}^*_{n-1}\|_{L_2} + \rho.\end{aligned}$$

By the strong convexity assumption A.2, we have

$$\frac{m}{2}\|\boldsymbol{x}_{n-1} - \boldsymbol{x}^*_{n-1}\|_2^2 \leq f_n(\boldsymbol{x}_{n-1}) - f_n(\boldsymbol{x}^*_{n-1}) \tag{13}$$

yielding

$$\mathbb{E}\|\boldsymbol{x}_{n-1} - \boldsymbol{x}^*_{n-1}\|_2^2 \leq \frac{2}{m}\left(\mathbb{E}[f_n(\boldsymbol{x}_{n-1})] - f_n(\boldsymbol{x}^*_{n-1})\right) \leq \frac{2\varepsilon}{m}.$$

Putting everything together we have

$$\mathbb{E}[d_n^2(0)] \leq \left(\sqrt{\frac{2\varepsilon}{m}} + \rho\right)^2. \tag{14}$$

Therefore, we have a bound

$$\mathbb{E}[f_n(\boldsymbol{x}_n)] - f_n(\boldsymbol{x}^*_n) \leq b\left(\sqrt{\frac{2\varepsilon}{m}} + \rho, K\right)$$

and we can set

$$K^* = \min\left\{ K \geq 1 \;\middle|\; b\left(\sqrt{\frac{2\varepsilon}{m}} + \rho, K\right) \leq \varepsilon \right\} \tag{15}$$

to ensure that

$$\mathbb{E}[f_n(\boldsymbol{x}_n)] - f_n(\boldsymbol{x}_n^*) \leq \varepsilon \quad \forall n \geq 1.$$

4 Estimating the Change in Minimizers

In practice, we do not know ρ, so we must construct an estimate $\hat{\rho}_n$ using the stochastic gradients $\{\boldsymbol{g}_n(\boldsymbol{x}, z_n(k))\}_{k=1}^{K_n}$. First, we construct estimates $\tilde{\rho}_i$ for the one step changes $\|\boldsymbol{x}_i^* - \boldsymbol{x}_{i-1}^*\|_2$ for each i. Next, we combine the one step estimates to construct an overall estimate $\hat{\rho}_n$ for ρ. As an intermediary step, we also look at a special case in which either

$$\|\boldsymbol{x}_{n+1}^* - \boldsymbol{x}_n^*\|_2 = \rho \tag{16}$$

or

$$\|\boldsymbol{x}_{n+1}^* - \boldsymbol{x}_n^*\|_{L_2} = \rho. \tag{17}$$

We show that for our estimate $\hat{\rho}_n$ and appropriately chosen sequences $\{t_n\}$, for all n large enough, $\hat{\rho}_n + t_n \geq \rho$ almost surely. With this property, analysis similar to that in Section 3 holds.

4.1 One Step Changes

We construct an estimate $\tilde{\rho}_i$ for the one-step changes $\|\boldsymbol{x}_i^* - \boldsymbol{x}_{i-1}^*\|$. As a consequence of the strong convexity of $f_i(\boldsymbol{x})$, we have the following lemma

Lemma 2. *It holds that* $\|\boldsymbol{x} - \boldsymbol{x}_i^*\|_2 \leq \frac{1}{m}\|\nabla f_i(\boldsymbol{x})\|_2 \quad \forall i \geq 1, \quad \forall \boldsymbol{x} \in \mathscr{X}$.

Proof. Since our functions $f_i(\boldsymbol{x})$ are convex, it holds that $\langle \nabla f_i(\boldsymbol{x}_i^*), \boldsymbol{x} - \boldsymbol{x}_i^* \rangle \geq 0 \forall \boldsymbol{x} \in \mathscr{X}$. By the strong monotonicity of the gradient, a consequence of strong convexity [25], it holds that $\langle \nabla f_i(\boldsymbol{x}) - \nabla f_i(\tilde{\boldsymbol{x}}), \boldsymbol{x} - \tilde{\boldsymbol{x}} \rangle \geq m\|\boldsymbol{x} - \tilde{\boldsymbol{x}}\|_2^2 \forall \boldsymbol{x}, \tilde{\boldsymbol{x}} \in \mathscr{X}$. Plugging in $\tilde{\boldsymbol{x}} = \boldsymbol{x}_i^*$ yields

$$\langle \nabla f_i(\boldsymbol{x}), \boldsymbol{x} - \boldsymbol{x}_i^* \rangle \geq m\|\boldsymbol{x} - \boldsymbol{x}_i^*\|_2^2.$$

Applying the Cauchy-Schwarz inequality yields the result. □

This in turn by way of the triangle inequality proves that

$$\begin{aligned} \|\boldsymbol{x}_i^* - \boldsymbol{x}_{i-1}^*\|_2 &\leq \|\boldsymbol{x}_i - \boldsymbol{x}_{i-1}\|_2 + \|\boldsymbol{x}_i - \boldsymbol{x}_i^*\|_2 + \|\boldsymbol{x}_{i-1} - \boldsymbol{x}_{i-1}^*\|_2 \\ &\leq \|\boldsymbol{x}_i - \boldsymbol{x}_{i-1}\|_2 + \frac{1}{m}\|\nabla f_i(\boldsymbol{x}_i)\|_2 + \frac{1}{m}\|\nabla f_{i-1}(\boldsymbol{x}_{i-1})\|_2. \end{aligned} \tag{18}$$

Motivated by this bound, we define the following estimate denoted the *direct estimate* by approximating the gradients

$$\tilde{\rho}_i \triangleq \|\boldsymbol{x}_i - \boldsymbol{x}_{i-1}\|_2 + \frac{1}{m}\|\hat{G}_i\|_2 + \frac{1}{m}\|\hat{G}_{i-1}\|_2 \tag{19}$$

where

$$\hat{G}_i = \frac{1}{K_i}\sum_{k=1}^{K_i} \boldsymbol{g}_i(\boldsymbol{x}_i, \boldsymbol{z}_i(k)).$$

4.2 Combining with Constant Change of Minimizers

As a special case, we look at combining the one step estimates when either eq. (16) or eq. (17) holds.

4.2.1 Euclidean Norm Condition

Under eq. (16), we construct an estimate by averaging the one step estimates

$$\hat{\rho}_n \triangleq \frac{1}{n-1}\sum_{i=2}^{n}\tilde{\rho}_i. \tag{20}$$

We want to show that for an appropriate sequence $\{t_n\}$, described in Theorem 1 and Theorem 3 below, and for all n large enough $\hat{\rho}_n + t_n \geq \rho$ almost surely under eq. (16) or eq. (17). The difficulty in actually proving this condition for eq. (19) is that when we compute

$$\hat{G}_i = \frac{1}{K_i}\sum_{k=1}^{K_i} \boldsymbol{g}_i\left(\boldsymbol{x}_i, \boldsymbol{z}_i(k)\right)$$

$\boldsymbol{x}_i$ and $\{\boldsymbol{z}_i(k)\}_{k=1}^{K_i}$ are dependent. To get around this issue, we consider performing a second independent draw of samples $\{\tilde{\boldsymbol{z}}_i(k)\}_{k=1}^{K_i}$. Note that we do not need to actually draw new independent samples; this is purely for the sake of analysis. We start from $\boldsymbol{x}_{i-1}$ and produce $\tilde{\boldsymbol{x}}_i$ using these new samples. For example, with SGD, we have

$$\begin{aligned}\boldsymbol{x}_i(k) &= \Pi_{\mathscr{X}}\left[\boldsymbol{x}_i(k-1) - \mu(k)\boldsymbol{g}_i\left(\boldsymbol{x}_i(k-1), \boldsymbol{z}_i(k)\right)\right]\\ \tilde{\boldsymbol{x}}_i(k) &= \Pi_{\mathscr{X}}\left[\tilde{\boldsymbol{x}}_i(k-1) - \mu(k)\boldsymbol{g}_i\left(\boldsymbol{x}_i(k-1), \tilde{\boldsymbol{z}}_i(k)\right)\right]\end{aligned} \tag{21}$$

for $k = 1, \ldots, K_i$ with $\boldsymbol{x}_i(0) = \tilde{\boldsymbol{x}}_i(0) = \boldsymbol{x}_{i-1}$. Then we copy the form of the direct estimate using $\tilde{\boldsymbol{x}}_i$ in place of $\boldsymbol{x}_i$ by defining

$$\tilde{\rho}_i^{(2)} = \|\tilde{\boldsymbol{x}}_i - \tilde{\boldsymbol{x}}_{i-1}\|_2 + \frac{1}{m}\|\tilde{G}_i\|_2 + \frac{1}{m}\|\tilde{G}_{i-1}\|_2 \tag{22}$$

with

$$\tilde{G}_i = \frac{1}{K_i}\sum_{k=1}^{K_i} \boldsymbol{g}_i\left(\tilde{\boldsymbol{x}}_i, \boldsymbol{z}_i(k)\right).$$

In this case, $\tilde{\boldsymbol{x}}_i$ and $\{\boldsymbol{z}_i(k)\}_{k=1}^{K_i}$ are independent, so $\mathbb{E}[\tilde{\rho}_i^{(2)}] \geq \rho$ by Lemma 2. Under eq. (16), using a dependent sub-Gaussian concentration inequality from [26] similar to Hoeffding's inequality, we then argue that $\hat{\rho}_n$ from eq. (20) is close to

$$\hat{\rho}_n^{(2)} = \frac{1}{n-1}\sum_{i=2}^{n}\tilde{\rho}_i^{(2)}$$

which in turn upper bounds ρ for all n large enough almost surely. Similarly, under eq. (17), we show that $\hat{\rho}_n^2$ from eq. (24) is close to $(\hat{\rho}_n^{(2)})^2$, which in turn upper bounds ρ^2 for all n large enough almost surely.

To proceed with our analysis, define the following σ-algebra

$$\mathscr{F}_i \triangleq \sigma\left(\bigcup_{j=1}^{i}\bigcup_{k=1}^{K_j} \boldsymbol{z}_j(k)\right)$$

and suppose that the following conditions hold:

B.1 Suppose there exist functions $C_i(K_i)$ such that $\mathbb{E}\left[\|\boldsymbol{x}_i - \tilde{\boldsymbol{x}}_i\|_2^2 \mid \mathscr{F}_{i-1}\right] \leq C_i^2(K_i)$.

B.2 Suppose that $\forall \boldsymbol{x}, \tilde{\boldsymbol{x}} \in \mathscr{X}$ it holds that

$$\mathbb{E}\left[\|\boldsymbol{g}_i\left(\boldsymbol{x}, \boldsymbol{z}_i\right) - \boldsymbol{g}_i\left(\tilde{\boldsymbol{x}}, \boldsymbol{z}_i\right)\|_2 \mid \mathscr{F}_{i-1}\right] \leq M\mathbb{E}\left[\|\boldsymbol{x} - \tilde{\boldsymbol{x}}\|_2 \mid \mathscr{F}_{i-1}\right]$$

and

$$\mathbb{E}\left[\|\boldsymbol{g}_i\left(\boldsymbol{x}, \boldsymbol{z}_i\right) - \nabla f_i(\boldsymbol{x})\|_2 \mid \mathscr{F}_{i-1}\right] \leq \sigma \;\; \forall \boldsymbol{x} \in \mathscr{X}.$$

B.3 Suppose the gradients are bounded in the sense that $\|\boldsymbol{g}_n\left(\boldsymbol{x}, \boldsymbol{z}\right)\|_2 \leq G \quad \forall \boldsymbol{x} \in \mathscr{X}, \boldsymbol{z} \in \mathscr{Z}$.

Assumption B.1 is a bound on the difference in how far apart two independent outputs of the optimization algorithm $\boldsymbol{x}_i$ and $\tilde{\boldsymbol{x}}_i$ starting from $\boldsymbol{x}_{i-1}$ are. assumption B.2 controls how the gradient grows for two points $\boldsymbol{x}$ and $\tilde{\boldsymbol{x}}$. Finally, assumption B.3 is reasonable if the space $\mathscr{Z}$ that contains the $\boldsymbol{z}_n$ has finite diameter. In this case, it holds that

$$\sup_{\boldsymbol{x}\in\mathscr{X},\boldsymbol{z}\in\mathscr{Z}} \|\boldsymbol{g}_n(\boldsymbol{x},\boldsymbol{z})\|_2 < +\infty.$$

In the following theorem, we establish that the direct estimate from eq. (20) (i.e., the Euclidean norm condition) upper bounds ρ from eq. (16) eventually.

Theorem 1. *Provided that assumptions B.1-B.3 hold and our sequence* $\{t_n\}$[1] *satisfies*

$$\sum_{n=2}^{\infty}\left(\exp\left\{-\frac{(n-1)t_n^2}{18diam^2(\mathscr{X})}\right\}+2\exp\left\{-\frac{m^2(n-1)t_n^2}{72G^2}\right\}\right)<+\infty$$

it holds that for all n large enough $\hat{\rho}_n + D_n + t_n \geq \rho$ *almost surely with*

$$D_n = \frac{1}{n-1}\left[\left(1+\frac{M}{m}\right)C_1+\sqrt{\frac{\sigma}{K_1}}+2\sum_{i=2}^{n-1}\left(\left(1+\frac{M}{m}\right)C_i+\sqrt{\frac{\sigma}{K_i}}\right)+\left(1+\frac{M}{m}\right)C_n+\sqrt{\frac{\sigma}{K_n}}\right]. \tag{23}$$

Proof. See Section 8 . □

From now on, for notational convenience, we absorb D_n into the t_n term and refer only to $\hat{\rho}_n + t_n$.

4.2.2 L_2 Norm Condition

Under eq. (17), we construct an estimate by averaging the squares of the one step estimates and taking a square root

$$\hat{\rho}_n \triangleq \sqrt{\frac{1}{n-1}\sum_{i=2}^{n}\tilde{\rho}_i^2}. \tag{24}$$

In the following theorem, we establish that the direct estimate from eq. (24) upper bounds ρ from eq. (17) eventually.

Theorem 2. *Provided that assumptions B.1-B.3 hold and our sequence* $\{t_n\}$ *satisfies*

$$\sum_{n=2}^{\infty}\left(\exp\left\{-\frac{(n-1)t_n^2}{18diam^2(\mathscr{X})}\right\}+2\exp\left\{-\frac{m^2(n-1)t_n^2}{72G^2}\right\}\right)<+\infty$$

it holds that for all n large enough $\sqrt{(\hat{\rho}_n)^2 + \tilde{D}_n + t_n} \geq \rho$ *almost surely with*
$\tilde{D}_n = 2diam(\mathscr{X})D_n$.

Proof. See Section 8. □

4.3 Combining with Bounded Changes of Minimizers

We examine estimating ρ in the case that eq. (2) holds. We denote the exact one step time changes by $\rho_i \triangleq \|\boldsymbol{x}_i^* - \boldsymbol{x}_{i-1}^*\|$. The simplest way to combine the estimates from eq. (19) would be to set

$$\hat{\rho}_n = \max\{\tilde{\rho}_2,\ldots,\tilde{\rho}_n\}.$$

For the sake of argument, suppose that $\tilde{\rho}_i = \rho_i + e_i$ with independent $e_i \sim \mathscr{N}(0,\sigma^2)$. Then it follows that [27]

$$\mathbb{E}[\hat{\rho}_n] \geq \mathbb{E}[\max\{e_2,\ldots,e_n\}].$$

For independent Gaussian random variables, it holds that $\mathbb{E}[\max\{e_2,\ldots,e_n\}] \to \infty$ as $n\to\infty$, and therefore this estimate goes to infinity as $n\to\infty$. We do produce an upper bound, but it increases to the trivial bound $\text{diam}^2(\mathscr{X})$. Next, we examine how to avoid this issue.

[1] Note that a choice of t_n that is no greater than $1/\sqrt{n-1}$ works here.

4.3.1 Euclidean Norm Condition

Suppose that the following conditions hold.

B.4 We have estimates $\hat{h}_W : \mathbb{R}^W \to \mathbb{R}$ that are non-decreasing in their arguments such that

$$\mathbb{E}[\hat{h}_W(\rho_j, \ldots, \rho_{j-W+1})] \geq \rho.$$

B.5 There exists absolute constants $\{a_i\}_{i=1}^W$ for any fixed $W \geq 1$ such that $\forall \boldsymbol{p}, \boldsymbol{q} \in \mathbb{R}^W_{\geq 0}$

$$|\hat{h}_W(p_1, \ldots, p_W) - \hat{h}_W(q_1, \ldots, q_W)| \leq \sum_{i=1}^{W} a_i |p_i - q_i|.$$

For example, if $\rho_i \overset{\text{iid}}{\sim} \text{Unif}[0, \rho]$, then

$$\hat{h}_W(\rho_i, \rho_{i+1}, \ldots, \rho_{i+W-1}) = \frac{W+1}{W} \max\{\rho_i, \rho_{i+1}, \ldots, \rho_{i+W-1}\} \tag{25}$$

is an estimate of ρ with the required properties with $a_i = 1 + \frac{1}{W}$. In this case, we compute the maximum in (25) over a sliding window and then average the maxima. This estimate will not blow up and will eventually upper bound ρ as we will see in Theorem 3 below.

Given an estimate satisfying assumptions B.4-B.5, we compute

$$\overline{\rho}^{(i)} = \hat{h}_W(\tilde{\rho}_i, \tilde{\rho}_{i-1}, \ldots, \tilde{\rho}_{i-W+1})$$

and produce an estimate $\hat{\rho}_n$ that is an average of these estimates

$$\hat{\rho}_n = \frac{1}{n-1} \sum_{i=2}^{n} \overline{\rho}^{(i)} = \frac{1}{n-1} \sum_{i=2}^{n} \hat{h}_{\min\{W, i-1\}}(\tilde{\rho}_i, \tilde{\rho}_{i-1}, \ldots, \tilde{\rho}_{\max\{i-W+1, 2\}}). \tag{26}$$

In the following theorem, we establish that the direct estimate from eq. (26) (i.e., the Euclidean norm condition) upper bounds ρ from eq. (2) eventually.

Theorem 3. *Provided that assumptions B.1-B.5 hold and our sequence $\{t_n\}$ is chosen such that*

$$\sum_{n=2}^{\infty} \left(\exp\left\{ -\frac{(n-W)^2 t_n^2}{18(n-1) diam^2(\mathscr{X}) \left(\sum_{j=1}^W a_j\right)^2} \right\} + 2\exp\left\{ -\frac{m^2 (n-W)^2 t_n^2}{72(n-1) G^2 \left(\sum_{j=1}^W a_j\right)^2} \right\} \right)$$

is finite, it holds that for all n large enough

$$\hat{\rho}_n + \left(\frac{n-1}{n-W} \sum_{j=1}^{W} a_j \right) D_n + t_n \geq \rho$$

with D_n from Theorem 1.

Proof. The proof in this case is similar to the proof for the equality assumption on ρ in eq. (16) and is provided in Section 8. □

As before, we will absorb $\left(\frac{n-1}{n-W} \sum_{j=1}^W a_j\right) D_n$ into t_n.

4.3.2 L_2 Norm Condition

Suppose that the following conditions hold, which are
analogs of assumptions B.4-B.5

B.6 We have estimates $\hat{h}_W : \mathbb{R}^W \to \mathbb{R}$ that are non-decreasing in their arguments such that

$$\mathbb{E}[\hat{h}_W(\rho_j^2, \ldots, \rho_{j-W+1}^2)] \geq \rho^2.$$

B.7 There exists absolute constants $\{a_i\}_{i=1}^W$ for any fixed $W \geq 1$ such that $\forall \boldsymbol{p}, \boldsymbol{q} \in \mathbb{R}_{\geq 0}^W$

$$|\hat{h}_W(p_1^2, \ldots, p_W^2) - \hat{h}_W(q_1^2, \ldots, q_W^2)| \leq \sum_{i=1}^W a_i |p_i^2 - q_i^2|.$$

For example, if $\rho_i \overset{\text{iid}}{\sim} \text{Unif}[0, \rho]$, then

$$\hat{h}_W\left(\rho_i^2, \rho_{i+1}^2, \ldots, \rho_{i+W-1}^2\right) = \frac{W+2}{W} \max\{\rho_i^2, \rho_{i+1}^2, \ldots, \rho_{i+W-1}^2\}$$

is an estimate of ρ with the required properties with $b_i = \frac{W+2}{W}$. In this case, we compute the max over a sliding window and then average the maximums. This estimate will not blow up but will eventually upper bound ρ as we will see later.

Given an estimate satisfying assumptions B.4-B.5, we compute

$$\bar{\rho}^{(i)} = \sqrt{\hat{h}_W(\tilde{\rho}_i^2, \tilde{\rho}_{i-1}^2, \ldots, \tilde{\rho}_{i-W+1}^2)}.$$

Under assumptions B.1-B.3 and assumptions B.6-B.7, we can then show that

$$\hat{\rho}_n = \sqrt{\frac{1}{n-1} \sum_{i=2}^n \left(\bar{\rho}^{(i)}\right)^2} \tag{27}$$

eventually upper bounds ρ.

Theorem 4. *Provided that assumptions B.1-B.3 and B.6-B.7 hold and our sequence $\{t_n\}$ is chosen such that the sum*

$$\sum_{n=2}^{\infty} \left(\exp\left\{ -\frac{(n-W)^2 t_n^2}{18(n-1)diam^2(\mathscr{X}) \left(\sum_{j=1}^W a_j\right)^2} \right\} + 2\exp\left\{ -\frac{m^2(n-W)^2 t_n^2}{72(n-1)G^2 \left(\sum_{j=1}^W a_j\right)^2} \right\} \right)$$

is finite, it holds that for all n large enough

$$\sqrt{(\hat{\rho}_n)^2 + \left(\frac{n-1}{n-W} \sum_{j=1}^W a_j\right) \tilde{D}_n + t_n} \geq \rho$$

with $\tilde{D}_n$ from Theorem 2.

Proof. The proof in this case is similar to the proof for the equality assumption on ρ in eq. (16) and is provided in Section 8. □

5 Tracking Analysis with Change in Minimizers Unknown

We now examine the case with ρ unknown. We extend the work of Section 3 using the estimate of ρ in Section 4. Our analysis depends on the following crucial assumptions:

C.1 For appropriate sequences $\{t_n\}$, for all n sufficiently large it holds that $\hat{\rho}_n + t_n \geq \rho$ almost surely.

C.2 The bound $b(d_0, K_n)$ defined in assumption A.4 factors as $b(d_0, K_n) = \alpha(K_n)d_0^2 + \beta(K_n)$.

We have demonstrated that assumption C.1 holds for the direct estimate of ρ.

In this section, we assume that either the constant change in minimizers condition, eq. (2), or the bounded change in minimizers condition, eq. (17), holds. Our analysis is not affected by which one is true. We use the following result, proved in Section 9, to derive rules to pick K_n when ρ is unknown:

Theorem 5. *Under assumptions C.1-C.2, with $K_n \geq K^*$ for all n large enough where K^* is defined in* (15) *we have*

$$\limsup_{n\to\infty} (\mathbb{E}[f_n(\boldsymbol{x}_n)] - f_n(\boldsymbol{x}_n^*)) \leq \varepsilon$$

almost surely

Proof. See Section 9. □

5.1 Update Past Mean Criterion Bounds

We first consider updating all past mean criterion bounds as we go. At time n, we plug-in $\hat{\rho}_{n-1} + t_{n-1}$ in place of ρ and follow the analysis of Section 3. Define

$$\hat{\varepsilon}_i^{(n)} \quad = \quad b\left(\sqrt{\frac{2}{m}\hat{\varepsilon}_{i-1}^{(n)}} + \hat{\rho}_{n-1} + t_{n-1}, K_i\right) \quad i = 1, \ldots, n.$$

If it holds that $\hat{\rho}_{n-1} + t_{n-1} \geq \rho$, then $\mathbb{E}[f_n(\boldsymbol{x}_n)] - f_n(\boldsymbol{x}_n^*) \leq \hat{\varepsilon}_n^{(i)}$ for $i = 1, \ldots, n$. assumption C.1 guarantees that this holds for all n large enough almost surely. We can thus set K_n equal to

$$K_n = \min\left\{K \;\middle|\; b\left(\sqrt{\frac{2}{m}\max\{\hat{\varepsilon}_{n-1}^{(n-1)}, \varepsilon\}} + \hat{\rho}_{n-1} + t_{n-1}, K\right) \leq \varepsilon\right\}$$

for all $n \geq 3$ to achieve mean criterion ε. The maximum in this definition ensures that when $\hat{\rho}_{n-1} + t_{n-1} \geq \rho$, $K_n \geq K^*$ with K^* from eq. (15). We can therefore apply Theorem 5.

5.2 Do Not Update Past Mean Criterion Bounds

Updating all past estimates of the mean criterion bounds from time 1 up to n imposes a computational and memory burden. Suppose that instead for all $n \geq 3$ we set

$$K_n = \min\left\{K \geq 1 \;\middle|\; b\left(\sqrt{\frac{2\varepsilon}{m}} + \hat{\rho}_{n-1} + t_{n-1}, K\right) \leq \varepsilon\right\}. \tag{28}$$

This is the same form as the choice in eq. (15) with $\hat{\rho}_{n-1} + t_{n-1}$ in place of ρ. Due to assumption C.1, for all n large enough it holds that $\hat{\rho}_n + t_n \geq \rho$ almost surely. Then by the monotonicity assumption in A.4, for all n large enough we would pick $K_n \geq K^*$ almost surely. We can therefore apply Theorem 5.

5.3 In High Probability Bounds

We discuss how to apply these results to the IHP criterion in Section 11.2.

6 Experiment

We apply our framework to a mean-squared vector estimation problem [2]. We fix the following signal model:

$$y_n = \boldsymbol{\eta}_n^\top \boldsymbol{w}_n + e_n.$$

Our goal is to estimate $\boldsymbol{\eta}_n$. We consider minimizing the following functions to estimate $\boldsymbol{\eta}_n$

$$f_n(\boldsymbol{x}) = \mathbb{E}_{\boldsymbol{z}_n \sim p_n}\left[\frac{1}{2}(y_n - \boldsymbol{x}^\top \boldsymbol{w}_n)^2\right]. \tag{29}$$

By simple algebraic manipulation, it holds that

$$f_n(\boldsymbol{x}) = \frac{1}{2}(\boldsymbol{x} - \boldsymbol{\eta}_n)^\top \mathbb{E}[\boldsymbol{w}_n \boldsymbol{w}_n^\top](\boldsymbol{x} - \boldsymbol{\eta}_n) + \frac{1}{2}\mathbb{E}[e_n^2]. \tag{30}$$

It is easy to see then that $\boldsymbol{x}_n^* = \boldsymbol{\eta}_n$. Set

$$\boldsymbol{z}_n \triangleq [\boldsymbol{w}_n^\top \; e_n]^\top \tag{31}$$

and define the stochastic gradients

$$\boldsymbol{g}_n(\boldsymbol{x}, \boldsymbol{z}_n) \triangleq -(y_n - \boldsymbol{x}^\top \boldsymbol{w}_n)\boldsymbol{w}_n$$

which satisfy the required condition in eq. (3). To find an approximation to $\boldsymbol{x}_n^*$, we apply SGD using the inverse step size averaging technique discussed in Section 10.

Let $\boldsymbol{w}_n \sim \mathcal{N}\left(\boldsymbol{0}, \frac{\sigma_w^2}{d}\boldsymbol{I}\right)$ where $\boldsymbol{w}_n \in \mathbb{R}^d$ and $e_n \sim \mathcal{N}(0, \sigma_e^2)$. We assume $\boldsymbol{\eta}_n$ is a deterministic sequence satisfying

$$\|\boldsymbol{\eta}_{n+1} - \boldsymbol{\eta}_n\|_2 \leq \rho. \tag{32}$$

Since $\boldsymbol{x}_n^* = \boldsymbol{\eta}_n$, the minimizer change condition in eq. (2) is satisfied. We use the ρ estimate in eq. (20). Note that since $\{\boldsymbol{\eta}_n\}$ is deterministic, we cannot apply a Kalman filter. Furthermore, we suppose that the collection of all $\boldsymbol{w}_n$, e_n, and $\boldsymbol{\eta}_n$ over all time instants are independent.

With this choice of model combined with the form of the functions in eq. (30), it is clear that the functions $f_n(\boldsymbol{x})$ are strongly convex with $m = \sigma_w^2/d$ satisfying assumption A.2. By applying the inequality $(a+b) \leq 2a^2 + 2b^2$, it follows that

$$\begin{aligned}
\mathbb{E}\|\boldsymbol{g}_n(\boldsymbol{x}, \boldsymbol{z}_n)\|_2^2 &= \mathbb{E}\|\boldsymbol{g}_n(\boldsymbol{x}^*, \boldsymbol{z}_n) + (\boldsymbol{g}_n(\boldsymbol{x}, \boldsymbol{z}_n) - \boldsymbol{g}_n(\boldsymbol{x}^*, \boldsymbol{z}_n))\|_2^2 \\
&\leq 2\mathbb{E}\|\boldsymbol{g}_n(\boldsymbol{x}^*, \boldsymbol{z}_n)\|_2^2 + 2\mathbb{E}\|\boldsymbol{g}_n(\boldsymbol{x}, \boldsymbol{z}_n) - \boldsymbol{g}_n(\boldsymbol{x}^*, \boldsymbol{z}_n)\|_2^2.
\end{aligned}$$

For the first term, we have

$$\begin{aligned}
\mathbb{E}\|\boldsymbol{g}_n(\boldsymbol{x}_n^*, \boldsymbol{z}_n)\|_2^2 &= \mathbb{E}\|e_n \boldsymbol{w}_n\|_2^2 \\
&= \mathbb{E}[e_n^2]\mathbb{E}\|\boldsymbol{w}_n\|_2^2 \\
&= \sigma_e^2 \sigma_w^2
\end{aligned}$$

and for the second term, we have

$$\begin{aligned}
\mathbb{E}\|\boldsymbol{g}_n(\boldsymbol{x}, \boldsymbol{z}_n) - \boldsymbol{g}_n(\boldsymbol{x}^*, \boldsymbol{z}_n)\|^2 &= \mathbb{E}\left[\|\boldsymbol{w}_n\|^2 (\boldsymbol{x} - \boldsymbol{x}^*)\boldsymbol{w}_n \boldsymbol{w}_n^\top (\boldsymbol{x} - \boldsymbol{x}^*)\right] \\
&\leq \mathbb{E}\left[\|\boldsymbol{w}_n\|^4\right] \|\boldsymbol{x} - \boldsymbol{x}^*)\|^2 \\
&\leq 3\sigma_w^4 \|\boldsymbol{x} - \boldsymbol{x}^*\|^2.
\end{aligned}$$

The last inequality follows since $\boldsymbol{w}_n$ is a centered Gaussian and therefore, it holds that $\mathbb{E}|\boldsymbol{w}_n|^4 \leq 3\mathbb{E}|\boldsymbol{w}_n|^2$ This implies that

$$\mathbb{E}\|\boldsymbol{g}_n(\boldsymbol{x}, \boldsymbol{z}_n)\|_2^2 \leq 2\sigma_e^2 \sigma_w^2 + 6\sigma_w^4 \|\boldsymbol{x} - \boldsymbol{x}^*\|_2^2.$$

[2] In [2], we apply the framework developed in this paper to a variety of machine learning problems similar to eq. (4) using real data.

Therefore, for assumption A.5, we can set

$$\begin{aligned} A &= 2\sigma_e^2\sigma_{\boldsymbol{w}}^2 \\ B &= 6\sigma_w^4. \end{aligned}$$

Putting it together, we have the parameters summarized in Table 1.

Parameter	Value
m	$\sigma_{\boldsymbol{w}}^2/d$
A	$2\sigma_e^2\sigma_{\boldsymbol{w}}^2$
B	$6\sigma_{\boldsymbol{w}}^4$

Table 1: Parameter Table

For this simulation, we choose $d = 2$, $\sigma_w^2 = 0.5$, $\sigma_e^2 = 0.5$, and $\rho = 1$.

6.1 Mean Tracking Criterion

First, we assume that ρ and all the parameters in Table 1 are known. We focus on the mean tracking criterion in eq. (5). Figure 1 shows the trade-off for the optimal ε versus K^* defined in eq. (15). Any pair (ε, K) located above this curve can be achieved, in the sense that by setting $K_n = K$, we achieve

$$\limsup_{n\to\infty} (\mathbb{E}[f_n(\boldsymbol{x}_n)] - f_n(\boldsymbol{x}_n^*)) \leq \varepsilon.$$

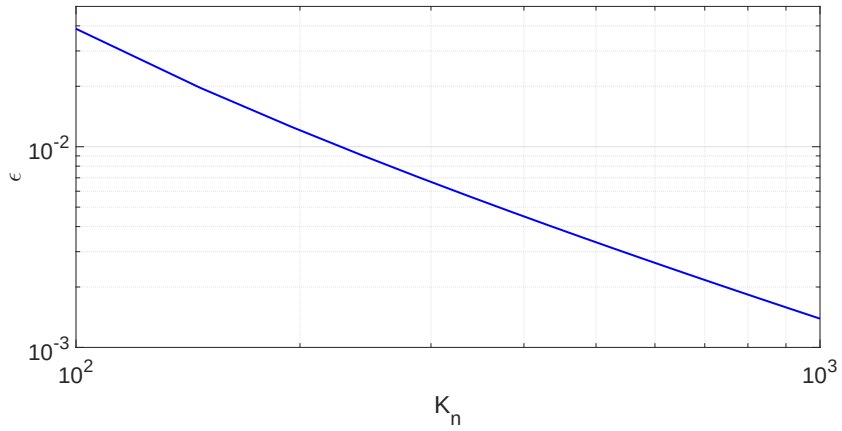

Figure 1: Mean tracking accuracy (ε) vs. Number of samples (K_n)

Next, we examine the case where ρ and the parameters in Table 1 are unknown. We estimate ρ using the techniques introduced in Section 4, specifically eq. (26), select K_n using the rule in eq. (28), and estimate the parameters using the techniques in Section 13. We target several different values of the mean tracking accuracy ε from eq. (5), including 0.001, 0.01, and 0.03. For the problem in this section, we can compute an estimate of the mean tracking accuracy to evaluate our methods. First, we have $f_n(\boldsymbol{x}_n^*) = \frac{1}{2}\sigma_e^2$. Second, for the sake of evaluation, we draw additional samples $\{\tilde{\boldsymbol{z}}_n(k)\}_{k=1}^{T_n} \overset{\text{iid}}{\sim} p_n$ and compute

$$\frac{1}{T_n}\frac{1}{2}\sum_{k=1}^{T_n}\left(\tilde{y}_n(k) - \boldsymbol{x}_n^\top \tilde{\boldsymbol{w}}_n(k)\right)^2$$

to estimate $f_n(\boldsymbol{x}_n)$. With these two pieces, we can estimate the tracking accuracy by computing

$$\frac{1}{T_n}\frac{1}{2}\sum_{k=1}^{T_n}\left(\tilde{y}_n(k)-\boldsymbol{x}_n^\top\tilde{\boldsymbol{w}}_n(k)\right)^2-\frac{1}{2}\sigma_e^2. \tag{33}$$

Table 2 shows an estimate of the actual achieved mean tacking accuracy for three different ε mean criterion targets averaged over $n=1$ to 100. In all cases, we meet our mean criterion target on average.

ε target	ε Estimate
0.001	0.0008 ± 0.0002
0.01	0.0073 ± 0.0012
0.03	0.022 ± 0.0022

Table 2: Estimate of mean criterion

Figure 2 shows the selected number of samples K_n for each mean tracking error target ε. In all cases we start from an insufficient number of samples $K_1 = K_2 = 50$. Due to the guarantees of 5, eventually we compensate for this initial bad choice to select K_n large enough. This process can be seen in Figure 2 by the spikes in K_n for small n to "catch up" to the correct K^*. For larger n, the choice of K_n settles down and does not vary greatly. Finally, as expected, for smaller choices of ε, K_n is larger.

Figure 3 shows the estimate of ρ. Our estimates of ρ upper bound the true value of $\rho = 1$ as desired. The initial spike in the estimates of ρ may be due to form of the $\hat{h}_W$ function in eq. (25) with $W = 4$. Before we have four one step estimates of ρ to plug in to $\hat{h}_4$, we use $\hat{h}_1, \hat{h}_2, \hat{h}_3$ per (25). The scaling factor in front of the maximum for these functions is $2, \frac{3}{2}, \frac{4}{3}$ before settling down to $\frac{5}{4}$. The larger scaling factors combined with the small number of one step ρ estimates results in an initial spike in the estimate of ρ. With more one step estimates of ρ, the overall estimate settles down. Finally, with a smaller mean tracking error target, we produce a tighter estimate of ρ.

Figure 4 shows the estimate $\hat{\varepsilon}_{i,n}$ of the mean tracking error achieved computed by updating the past. As mentioned above, we have an insufficient initial choice of K_1 and K_2, which causes initial spikes in the estimate of mean tracking error. Our choice of K_n drive these mean tracking error estimates below their target values of ε.

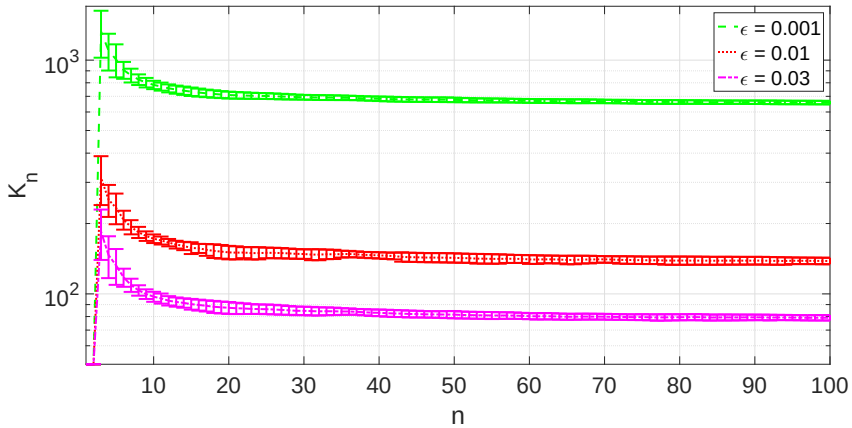

Figure 2: Number of samples (K_n)

6.2 IHP Tracking Criterion

Section 14.1contains a simulation for the IHP criterion.

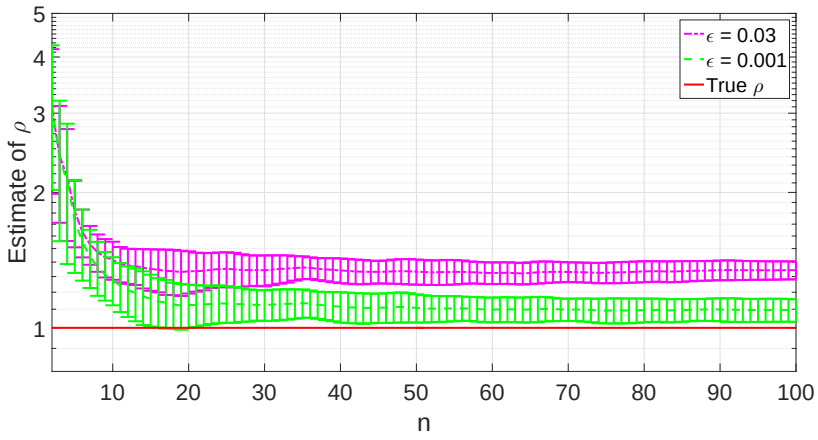

Figure 3: Estimate of change in minimizers (ρ)

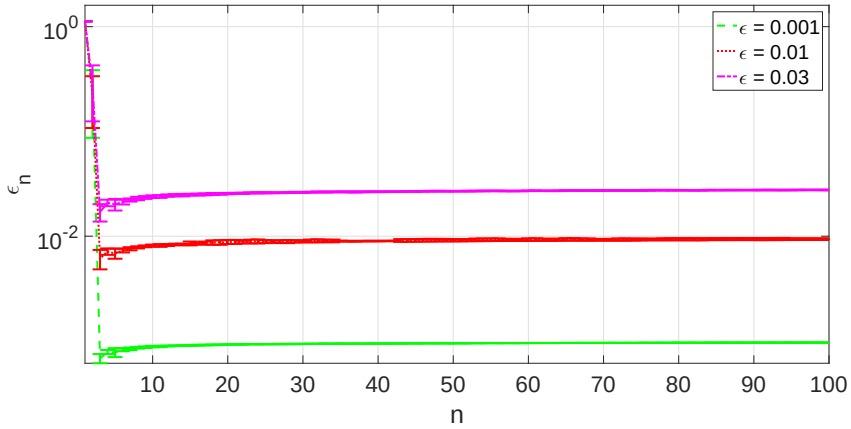

Figure 4: Estimate of mean tracking accuracy

7 Conclusion

We have developed a framework for sequentially solving stochastic optimization problems to within a target accuracy in a mean sense and in a high probability sense. An estimate of the change in the minimizers, combined with properties of the chosen optimization algorithm, is used to select the number of stochastic gradients needed to meet a given criterion. Simulations of our method demonstrate that our approach works well.

References

[1] C. Wilson, V. Veeravalli, and A. Nedić, "Dynamic stochastic optimization," in *CDC*, Dec. 2014.

[2] C. Wilson and V. Veeravalli, "Adaptive sequential optimization with applications to machine learning," in *ICASSP*, Mar. 2016.

[3] N. Cesa-Bianchi and G. Lugosi, *Prediction, Learning, and Games*. New York, N.Y., USA: Cambridge University Press, 2006.

[4] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning and stochastic optimization," *Journal of Machine Learning Research*, vol. 12, pp. 2121–2159, 2011.

[5] J. Duchi and Y. Singer, "Efficient online and batch learning using forward backward splitting," *Journal of Machine Learning Research*, vol. 10, pp. 2899–2934, 2009.

[6] E. Hazan, A. Agarwal, and S. Kale, "Logarithmic regret algorithms for online convex optimization," *Machine Learning*, vol. 69, pp. 169–192, 2007.

[7] E. H. P. Bartlett and A. Rakhlin, "Adaptive online gradient descent," in *Advances in Neural Information Processing Systems*, J. Platt, D. Koller, Y. Singer, and S. Roweis, Eds. Cambridge, MA, USA: MIT Press, 2008, vol. 20, pp. 65–72.

[8] S. Shalev-Shwartz and S. Kakade, "Mind the duality gap: Logarithmic regret algorithms for online optimization," in *Advances in Neural Information Processing Systems*, D. Koller, D.Schuurmans, Y. Bengio, and L. Bottou, Eds., vol. 21, MIT Press, 2009, pp. 1457–1464.

[9] S. Shalev-Shwartz and Y. Singer, "Convex repeated games and Fenchel duality," in *Advances in Neural Information Processing Systems*, vol. 19, MIT Press, Cambridge, MA, USA, 2006, pp. 1265–1271.

[10] ——, "Logarithmic regret algorithms for strongly convex repeated games," 2007, in The Hebrew University.

[11] L. Xiao, "Dual averaging methods for regularized stochastic learning and online optimization," Microsoft, Tech. Rep., March 2010, no. MSR-TR-2010-23.

[12] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *Proceedings of the 20th International Conference on Machine Learning (ICML)*, 2003, pp. 928–936.

[13] A. Rakhlin and K. Sridharan, "Online Learning with Predictable Sequences," *ArXiv*, vol. abs/1208.3728, Aug. 2012.

[14] C.-K. Chiang, T. Yang, C.-J. Lee, M. Mahdavi, C.-J. Lu, R. Jin, and S. Zhu, "Online optimization with gradual variations," in *COLT*, 2012.

[15] A. Dontchev and R. Rockafellar, *Implicit Functions and Solution Mappings: A View from Variational Analysis*. New York, New York: Springer, 2009.

[16] N. Takahashi, I. Yamada, and A. Sayed, "Diffusion least-mean squares with adaptive combiners: Formulation and performance analysis," *IEEE Transactions on Signal Processing*, vol. 58, no. 9, pp. 4795–4810, 2010.

[17] I. Yamada and N. Ogura, "Adaptive projected subgradient method for asymptotic minimization of sequence of nonnegative convex functions," *Taylor & Francis Group*, vol. 25, no. (7-8), pp. 593–617, 2005.

[18] K. Slavakis, I. Yamada, and N. Ogura, "The adaptive projected subgradient method over the fixed point set of strongly attracting nonexpansive mappings," *Numerical functional analysis and optimization*, vol. 27, no. (7-8), pp. 905–930, 2006.

[19] H. Kushner and J. Yang, "Analysis of adaptive step size sa algorithms for parameter tracking," in *Decision and Control, 1994., Proceedings of the 33rd IEEE Conference on*, vol. 1, 1994, pp. 730–737 vol.1.

[20] V. Solo and X. Kong, *Adaptive Signal Processing Algorithms: Stability and Performance*. Englewood Cliffs, NJ: Prentice Hall, 1995.

[21] A. Sayed, *Adaptive Filters*. Hoboken, New Jersey, USA: Wiley & Sons, Inc., 2008.

[22] A. Doucet and A. M. Johansen. (2008, Dec.) A Tutorial on Particle Filtering and Smoothing: Fifteen years later.

[23] A. Müller, "Integral Probability Metrics and Their Generating Classes of Functions," *Advances in Applied Probability*, vol. 29, no. 2, pp. 429+, Jun. 1997.

[24] W. Rudin, *Principles of mathematical analysis*, ser. International series in pure and applied mathematics.

[25] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY, USA: Cambridge University Press, 2004.

[26] R. Antonini and Y. Kozachenko, "A note on the asymptotic behavior of sequences of generalized subgaussian random vectors," *Random Op. and Stoch. Equ.*, vol. 13, pp. 39–52, 2005.

[27] R. Durrett, *Probability: Theory and examples*. Cambridge University Press, 2011.

[28] S. Boucheron, G. Lugosi, and P. Massart, *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.

[29] J. Kennan, "Uniqueness of positive fixed points for increasing concave functions on rn: An elementary result," *Review of Economic Dynamics*, vol. 4, pp. 893–899, 2001.

[30] A. Granas and J. Dugundji, *Fixed Point Theory*. Springer-Verlag, 2003.

[31] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro, "Stochastic approximation approach to stochastic programming," *SIAM Journal on Optimization*, vol. 19, pp. 1574–1609, 2009.

[32] F. Bach and E. Moulines, "Non-Asymptotic Analysis of Stochastic Approximation Algorithms for Machine Learning," in *Advances in Neural Information Processing Systems (NIPS)*, Spain, 2011.

[33] D. Bertsekas, *Nonlinear Programming*. Athena Scientific, 1999.

[34] A. Nedic and S. Lee, "Analysis of mirror descent for strongly convex functions," *ArXiV*, 2013.

[35] Y. Nesterov, *Introductory Lectures on Convex Optimization: A Basic Course*. Norwell, Massachusetts, USA: Kluwer Academic Publishers, 2004.

[36] S. Haykin, *Adaptive Filter Theory*. Springer, 2002.

8 Proofs for Estimates of Change in Minimizers

For our analysis of minimizer change estimation, we need to introduce a few results for sub-Gaussian random variables including the following key technical lemma from [26]. This lemma controls the concentration of sums of random variables that are sub-Gaussian conditioned on a particular filtration $\{\mathscr{F}_i\}_{i=0}^n$. Such a collection of random variables is referred to as a *sub-Gaussian martingale sequence*.

Lemma 3 (Theorem 7.5 of [26])**.** *Suppose we have a collection of random variables $\{V_i\}_{i=1}^n$ and a filtration $\{\mathscr{F}_i\}_{i=0}^n$ such that for each random variable V_i it holds that*

1. $\mathbb{E}\left[\exp\left\{s\left(V_i - \mathbb{E}\left[V_i \,\middle|\, \mathscr{F}_{i-1}\right]\right)\right\} \,\middle|\, \mathscr{F}_{i-1}\right] \leq e^{\frac{1}{2}\sigma_i^2 s^2}$ *with σ_i^2 a constant.*
2. *V_i is $\mathscr{F}_i$-measurable.*

Then for every $\boldsymbol{a} \in \mathbb{R}^n$ it holds that

$$\mathbb{P}\left\{\sum_{i=1}^n a_i V_i > \sum_{i=1}^n a_i \mathbb{E}\left[V_i \,\middle|\, \mathscr{F}_{i-1}\right] + t\right\} \leq \exp\left\{-\frac{t^2}{2v}\right\}$$

with $v = \sum_{i=1}^n \sigma_i^2 a_i^2$. The other tail is similarly bounded.

If we can upper bound the conditional expectations $\mathbb{E}\left[V_i \,\middle|\, \mathscr{F}_{i-1}\right] \leq C_i$ by $\mathscr{F}_{i-1}$-measurable random variables C_i, then we have

$$\mathbb{P}\left\{\sum_{i=1}^n a_i V_i > \sum_{i=1}^n a_i C_i + t\right\} \leq \mathbb{P}\left\{\sum_{i=1}^n a_i V_i > \sum_{i=1}^n a_i \mathbb{E}\left[V_i \,\middle|\, \mathscr{F}_{i-1}\right] + t\right\} \leq \exp\left\{-\frac{t^2}{2\nu}\right\}.$$

For our analysis, we generally cannot compute $\mathbb{E}\left[V_i \,\middle|\, \mathscr{F}_{i-1}\right]$, but we can find "nice" C_i.

To find σ_i^2 for use in Lemma 3, we employ the following conditional version of Hoeffding's Lemma.

Lemma 4 (Conditional Hoeffding's Lemma)**.** *If a random variable V and a sigma algebra $\mathscr{F}$ satisfy $a \leq V \leq b$ and $\mathbb{E}[V|\mathscr{F}] = 0$, then*

$$\mathbb{E}\left[e^{sV} \mid \mathscr{F}\right] \leq \exp\left\{\frac{1}{8}(b-a)^2 s^2\right\}.$$

Proof. Follows from the standard proof of Hoeffding's Lemma from [28]. □

Using these tools, we can analyze averages of the direct estimate. We focus on the proof of Theorem 1 and Theorem 3 as the proofs of Theorem 2 and Theorem 4 are simple extensions.

8.1 Euclidean Norm Condition

As a reminder, we consider running our optimization algorithm used to generate $\boldsymbol{x}_i$ again using independent samples $\{\tilde{\boldsymbol{z}}_i(k)\}_{k=1}^{K_i}$ to yield a second approximate minimizer $\tilde{\boldsymbol{x}}_i$. For SGD, the process to do this is summarized in eq. (21). We connect $\tilde{\rho}_i$ to $\tilde{\rho}_i^{(2)}$ with $\tilde{\rho}_i^{(2)}$ defined in eq. (22).

Proof of Theorem 1. To proceed, we compare the three single step estimates:

1. $\tilde{\rho}_i = \|\boldsymbol{x}_i - \boldsymbol{x}_{i-1}\|_2 + \frac{1}{m}\|G_i\|_2 + \frac{1}{m}\|G_{i-1}\|_2$
2. $\tilde{\rho}_i^{(2)} = \|\tilde{\boldsymbol{x}}_i - \tilde{\boldsymbol{x}}_{i-1}\|_2 + \frac{1}{m}\|\tilde{G}_i\|_2 + \frac{1}{m}\|\tilde{G}_{i-1}\|_2$
3. $\tilde{\rho}_i^{(3)} = \|\tilde{\boldsymbol{x}}_i - \tilde{\boldsymbol{x}}_{i-1}\|_2 + \frac{1}{m}\|\nabla f_i(\tilde{\boldsymbol{x}}_i)\|_2 + \frac{1}{m}\|\nabla f_{i-1}(\tilde{\boldsymbol{x}}_{i-1})\|_2$

where

$$\hat{G}_i = \frac{1}{K_i}\sum_{k=1}^{K_i} \nabla_{\boldsymbol{x}} \ell(\boldsymbol{x}_i, \boldsymbol{z}_i(k))$$

and

$$\tilde{G}_i = \frac{1}{K_i}\sum_{k=1}^{K_i} \nabla_{\boldsymbol{x}} \ell(\tilde{\boldsymbol{x}}_i, \boldsymbol{z}_i(k)).$$

Define $\hat{\rho}_n^{(2)}$ and $\hat{\rho}_n^{(3)}$ analogously to $\hat{\rho}_n$ as an average of the relevant single step estimates.

Using the triangle inequality and the reverse triangle inequality, it holds that

$$\begin{aligned}
|\hat{\rho}_n - \hat{\rho}_n^{(3)}| &= |\hat{\rho}_n - \hat{\rho}_n^{(2)} + \hat{\rho}_n^{(2)} - \hat{\rho}_n^{(3)}| \\
&\leq |\hat{\rho}_n - \hat{\rho}_n^{(2)}| + |\hat{\rho}_n^{(2)} - \hat{\rho}_n^{(3)}| \\
&\leq \frac{1}{n-1}\sum_{i=2}^{n}\left(\|\boldsymbol{x}_i - \tilde{\boldsymbol{x}}_i\|_2 + \|\boldsymbol{x}_{i-1} - \tilde{\boldsymbol{x}}_{i-1}\|_2 + \frac{1}{m}\|\hat{G}_i - \tilde{G}_i\|_2 + \frac{1}{m}\|\hat{G}_{i-1} - \tilde{G}_{i-1}\|_2\right) \\
&\quad + \frac{1}{n-1}\sum_{i=2}^{n}\left(\frac{1}{m}\|\tilde{G}_i - \nabla f_i(\tilde{\boldsymbol{x}}_i)\|_2 + \frac{1}{m}\|\tilde{G}_{i-1} - \nabla f_{i-1}(\tilde{\boldsymbol{x}}_{i-1})\|_2\right) \\
&\leq \frac{1}{n-1}\left(\|\boldsymbol{x}_1 - \tilde{\boldsymbol{x}}_1\|_2 + 2\sum_{i=2}^{n-1}\|\boldsymbol{x}_i - \tilde{\boldsymbol{x}}_i\|_2 + \|\boldsymbol{x}_n - \tilde{\boldsymbol{x}}_n\|_2\right) \\
&\quad + \frac{1}{m(n-1)}\left(\|\hat{G}_1 - \tilde{G}_1\|_2 + 2\sum_{i=2}^{n-1}\|\hat{G}_i - \tilde{G}_i\|_2 + \|\hat{G}_n - \tilde{G}_n\|_2\right) \\
&\quad + \frac{1}{m(n-1)}\left(\|\tilde{G}_1 - \nabla f_1(\tilde{\boldsymbol{x}}_1)\|_2 + 2\sum_{i=2}^{n-1}\|\tilde{G}_i - \nabla f_i(\tilde{\boldsymbol{x}}_i)\|_2 + \|\tilde{G}_n - \nabla f_n(\tilde{\boldsymbol{x}}_n)\|_2\right).
\end{aligned}$$

Define

$$U_n = \frac{1}{n-1}\left(\|\boldsymbol{x}_1 - \tilde{\boldsymbol{x}}_1\|_2 + 2\sum_{i=2}^{n-1}\|\boldsymbol{x}_i - \tilde{\boldsymbol{x}}_i\|_2 + \|\boldsymbol{x}_n - \tilde{\boldsymbol{x}}_n\|_2\right)$$

and

$$V_n = \frac{1}{m(n-1)}\left(\|\hat{G}_1 - \tilde{G}_1\|_2 + 2\sum_{i=2}^{n-1}\|\hat{G}_i - \tilde{G}_i\|_2 + \|\hat{G}_n - \tilde{G}_n\|_2\right)$$

and

$$W_n = \frac{1}{m(n-1)}\left(\|\tilde{G}_1 - \nabla f_1(\tilde{\boldsymbol{x}}_1)\|_2 + 2\sum_{i=2}^{n-1}\|\tilde{G}_i - \nabla f_i(\tilde{\boldsymbol{x}}_i)\|_2 + \|\tilde{G}_n - \nabla f_n(\tilde{\boldsymbol{x}}_n)\|_2\right).$$

Then it holds that

$$|\hat{\rho}_n - \hat{\rho}_n^{(3)}| \leq U_n + V_n + W_n.$$

Now, we look at bounding $\mathbb{E}[\|\boldsymbol{x}_i - \boldsymbol{x}_{i-1}\|_2 \mid \mathscr{F}_{i-1}]$, $\mathbb{E}[\|\hat{G}_i - \tilde{G}_i\|_2 \mid \mathscr{F}_{i-1}]$, and $\mathbb{E}[\|\tilde{G}_i - \nabla f_i(\tilde{\boldsymbol{x}})_i\|_2 \mid \mathscr{F}_{i-1}]$. First, by assumption B.1, it holds that

$$\mathbb{E}[\|\boldsymbol{x}_i - \tilde{\boldsymbol{x}}_i\|_2 \mid \mathscr{F}_{i-1}] \leq C_i.$$

Second, it holds that

$$\begin{aligned}
\mathbb{E}\left[\|\hat{G}_i - \tilde{G}_i\|_2 \mid \mathscr{F}_{i-1}\right] &\leq \frac{1}{K_i}\sum_{k=1}^{K_i}\mathbb{E}[\|\boldsymbol{g}_i(\boldsymbol{x}_i, \boldsymbol{z}_i(k)) - \boldsymbol{g}_i(\tilde{\boldsymbol{x}}_i, \boldsymbol{z}_i(k))\|_2 \mid \mathscr{F}_{i-1}] \\
&= \mathbb{E}[\|\boldsymbol{g}_i(\boldsymbol{x}_i, \boldsymbol{z}_i(1)) - \boldsymbol{g}_i(\tilde{\boldsymbol{x}}_i, \boldsymbol{z}_i(1))\|_2 \mid \mathscr{F}_{i-1}] \\
&\leq M\mathbb{E}[\|\boldsymbol{x}_i - \tilde{\boldsymbol{x}}_i\|_2 \mid \mathscr{F}_{i-1}] \\
&\leq MC_i.
\end{aligned}$$

Third, it holds that

$$\begin{aligned}\mathbb{E}\left[\|\tilde{G}_i - \nabla f_i(\tilde{\boldsymbol{x}}_i)\|_2 \mid \mathscr{F}_{i-1}\right] &\leq \left(\mathbb{E}\left[\left\|\frac{1}{K_i}\sum_{k=1}^{K_i}\left(\boldsymbol{g}_i\left(\tilde{\boldsymbol{x}}_i, \boldsymbol{z}_i(k)\right) - \nabla f_i(\tilde{\boldsymbol{x}}_i)\right)\right\|_2^2 \,\middle|\, \mathscr{F}_{i-1}\right]\right)^{1/2} \\ &\leq \left(\mathbb{E}\left[\frac{1}{K_i^2}\sum_{k=1}^{K_i}\left\|\boldsymbol{g}_i\left(\tilde{\boldsymbol{x}}_i, \boldsymbol{z}_i(k)\right) - \nabla f_i(\tilde{\boldsymbol{x}}_i)\right\|_2^2 \,\middle|\, \mathscr{F}_{i-1}\right]\right)^{1/2} \\ &\leq \left(\frac{\sigma}{K_i}\right)^{1/2}.\end{aligned}$$

The resulting bounds on the expectation of U_n, V_n, and W_n denoted $\bar{U}_n$, $\bar{V}_n$, and $\bar{W}_n$ are as follows:

1. $\bar{U}_n = \frac{1}{n-1}\left(C_1 + 2\sum_{i=2}^{n-1} C_i + C_n\right)$
2. $\bar{V}_n = \frac{M}{m(n-1)}\left(C_1 + 2\sum_{i=2}^{n-1} C_i + C_n\right)$
3. $\bar{W}_n = \frac{1}{m(n-1)}\left(\left(\frac{\sigma}{K_1}\right)^{1/2} + 2\sum_{i=2}^{n-1}\left(\frac{\sigma}{K_1}\right)^{1/2} + \left(\frac{\sigma}{K_n}\right)^{1/2}\right)$.

Then it holds that

$$\begin{aligned}\mathbb{P}\left\{|\hat{\rho}_n - \hat{\rho}_n^{(3)}| > (\bar{U}_n + \bar{V}_n + \bar{W}_n) + t_n\right\} &\leq \mathbb{P}\left\{U_n + V_n + W_n > (\bar{U}_n + \bar{V}_n + \bar{W}_n) + t_n\right\} \\ &\leq \mathbb{P}\left\{U_n > \bar{U}_n + \frac{1}{3}t_n\right\} + \mathbb{P}\left\{V_n > \bar{V}_n + \frac{1}{3}t_n\right\} + \mathbb{P}\left\{W_n > \bar{W}_n + \frac{1}{3}t_n\right\}.\end{aligned}$$

Now, we bound each of these three probabilities using Lemma 3. First, we have

$$0 \leq \|\boldsymbol{x}_i - \tilde{\boldsymbol{x}}_i\|_2 \leq \mathrm{diam}(\mathscr{X})$$

so applying Lemma 4 and Lemma 3 with $\sigma_i^2 = \frac{1}{4}\mathrm{diam}^2(\mathscr{X})$ and

$$\begin{aligned}a_1 &= a_n = \frac{1}{n-1} \\ a_2 &= \cdots = a_{n-2} = \frac{2}{n-1}\end{aligned}$$

yields

$$\begin{aligned}v_U &= \frac{1}{4}\mathrm{diam}^2(\mathscr{X})\sum_{i=1}^{n} a_i^2 \\ &= \frac{1}{4}\mathrm{diam}^2(\mathscr{X})\left(\frac{1}{n-1}\right)^2 + \sum_{i=2}^{n-1}\left(\frac{2}{n-1}\right)^2 + \left(\frac{1}{n-1}\right)^2 \\ &\leq \frac{1}{n-1}\mathrm{diam}^2(\mathscr{X}). \qquad (34)\end{aligned}$$

Therefore, it holds that

$$\begin{aligned}\mathbb{P}\left\{U_n > \bar{U}_n + \frac{1}{3}t_n\right\} &\leq \exp\left\{-\frac{(t_n/3)^2}{2v_U}\right\} \\ &= \exp\left\{-\frac{(n-1)t_n^2}{18\mathrm{diam}^2(\mathscr{X})}\right\}.\end{aligned}$$

Since

$$0 \leq \|\hat{G}_i - \tilde{G}_i\|_2 \leq 2G$$

and

$$0 \leq \|\tilde{G}_i - \nabla f_i(\tilde{\boldsymbol{x}}_i)\|_2 \leq 2G$$

we can apply Lemma 4 and Lemma 3 to V_n and W_n to yield

$$\mathbb{P}\left\{V_n > \bar{V}_n + \frac{1}{3}t_n\right\} \leq \exp\left\{-\frac{(t_n/3)^2}{2v_V}\right\} = \exp\left\{-\frac{m^2(n-1)t_n^2}{72G^2}\right\}$$

and similarly

$$\mathbb{P}\left\{W_n > \bar{W}_n + \frac{1}{3}t_n\right\} \leq \exp\left\{-\frac{m^2(n-1)t_n^2}{72G^2}\right\}.$$

Define

$$D_n \quad = \quad \bar{U}_n + \bar{V}_n + \bar{W}_n$$

which is the definition in eq. (23). It follows that

$$\mathbb{P}\left\{\hat{\rho}_n < \hat{\rho}_n^{(3)} - D_n - t_n\right\} \leq \exp\left\{-\frac{(n-1)t_n^2}{18\text{diam}^2(\mathscr{X})}\right\} + 2\exp\left\{-\frac{m^2(n-1)t_n^2}{72G^2}\right\}.$$

Then it follows that

$$\begin{aligned}
&\sum_{n=2}^{\infty} \mathbb{P}\left\{\hat{\rho}_n < \hat{\rho}_n^{(3)} - D_n - t_n\right\} \\
&\qquad \leq \sum_{n=2}^{\infty}\left(\exp\left\{-\frac{(n-1)t_n^2}{18\text{diam}^2(\mathscr{X})}\right\}\right. \\
&\qquad\qquad \left. +2\exp\left\{-\frac{m^2(n-1)t_n^2}{72G^2}\right\}\right) \\
&\qquad < +\infty.
\end{aligned}$$

Therefore, by the Borel-Cantelli Lemma, for all n large enough it holds that

$$\hat{\rho}_n + D_n + t_n \geq \hat{\rho}_n^{(3)}$$

almost surely. Finally, by 2, it holds that $\hat{\rho}_n^{(3)} \geq \rho$, which proves the result. □

Looking at the form of D_i, it follows that in this case

$$D_n = \mathscr{O}\left(\frac{1}{n-1}\sum_{i=1}^{n}\frac{1}{\sqrt{K_i}}\right).$$

In the case where $K_i = K^*$, this implies that

$$D_n = \mathscr{O}\left(\frac{1}{\sqrt{K^*}}\right).$$

In Section 12, we finish the proofs for the inequality condition on ρ and the L_2 condition for equality and inequality. The arguments are similar to those employed here.

9 Proofs for Analysis with Change in Minimizers Unknown

We prove a general result showing that for any choice of K_n such that $K_n \geq K^*$ for all n large enough with K^* from eq. (15), the mean criterion is controlled in the sense that

$$\limsup_{n\to\infty} (\mathbb{E}[f_n(\boldsymbol{x}_n)] - f_n(\boldsymbol{x}_n^*)) \leq \varepsilon.$$

Consider the function

$$\phi_K(v) = \alpha(K)\left(\sqrt{\frac{2}{m}v} + \rho\right)^2 + \beta(K) = b\left(\sqrt{\frac{2}{m}v} + \rho, K\right) \tag{35}$$

from assumption C.2. Note that as a function of v, $\phi_K(v)$ is clearly increasing and strictly concave. If we select K^* defined in eq. (15), then by definition it holds that

$$\phi_{K^*}(\varepsilon) \leq \varepsilon. \tag{36}$$

First, we study fixed points of the function $\phi_{K^*}(v)$. We need Theorem 3.3 of [29] to proceed.

Lemma 5 (Theorem 3.3 of [29])**.** *Suppose that f is an increasing and strictly concave mapping from $\mathbb{R}$ to $\mathbb{R}$ such that $f(0) \geq 0$ and there exist points $0 < a < b$ such that $f(a) > a$ and $f(b) < b$. Then f has unique positive fixed point.*

Proof. See [29] for the proof. □

We consider the fixed points of the function $\phi_{K^*,\rho}(v) + \delta$ with $\delta \geq 0$. We add the term δ for reasons that will become clear later in the proof of Theorem 5.

Lemma 6. *Provided that $\alpha(K) > 0$ for all $K > 0$, $\rho > 0$, and $\delta \geq 0$, the function $\phi_{K^*,\rho}(v) + \delta$ has a unique positive fixed point $\bar{v}_\delta$ with the following properties:*

1. $\bar{v}_0 = \phi_{K^*,\rho}(\bar{v}_0) \leq \varepsilon$.
2. $\phi'_{K^*,\rho}(\bar{v}_\delta) < 1$.
3. *$\bar{v}_\delta$ is non-decreasing in δ and*
$$\lim_{\delta \searrow 0} \bar{v}_\delta = \bar{v}_0.$$

Proof. We have

$$\phi_{K^*}(0) + \delta = \alpha(K^*)\rho^2 + \beta(K^*) + \delta > 0.$$

Since

$$\lim_{v\to 0} (\phi_{K^*}(v) + \delta) = \phi_{K^*}(0) + \delta$$

and $\phi_{K^*}(0) > 0$, for all $\delta \geq 0$, there exists a positive a sufficiently small that

$$\phi_{K^*}(a) + \delta > a.$$

Next, expanding $\phi_K(v)$ yields

$$\phi_K(v) = \frac{2}{m}\alpha(K)v + 2\alpha(K)\rho\sqrt{\frac{2}{m}}\sqrt{v} + \alpha(K)\rho^2 + \beta(K).$$

Since $\phi_{K^*}(\varepsilon) \leq \varepsilon$, we obviously must have $\frac{2}{m}\alpha(K^*) \leq 1$. Suppose that

$$\frac{2}{m}\alpha(K^*) = 1.$$

Then it holds that

$$\phi_{K^*}(\varepsilon) = \varepsilon + \sqrt{2m}\rho\sqrt{\varepsilon} + \frac{m}{2}\rho^2 + \beta(K) > \varepsilon.$$

This contradicts eq. (36), so it holds that

$$\frac{2}{m}\alpha(K^*) < 1.$$

It is thus readily apparent that

$$v - (\phi_{K^*}(v) + \delta) \to \infty$$

as $v \to \infty$. Therefore, there exists a point $b > a$ such that

$$\phi_{K^*}(b) + \delta < b.$$

In addition, it is easy to check that $\phi_{K^*}(v) + \delta$ is increasing and strictly concave. Therefore, we can apply Lemma 5 from [29] to conclude that there exists a unique, positive fixed point $\bar{v}_\delta$ of $\phi_{K^*}(v) + \delta$.

Next, suppose that $\phi'_{K^*}(\bar{v}_\delta) > 1$. Then by continuity for $v > \bar{v}_\delta$ sufficiently close to $\bar{v}_\delta$, we have

$$\phi_{K^*}(v) + \delta > v.$$

However, we know that as $v \to \infty$, it holds that $v - (\phi_{K^*}(v) + \delta) \to \infty$. By the Intermediate Value Theorem, this implies that there is another fixed point on $[v, b]$. This is a contradiction, since $\bar{v}_\delta$ is the unique, positive fixed point. Therefore, it holds that $\phi'_{K^*}(\bar{v}_\delta) \leq 1$. Now, suppose that $\phi'_{K^*}(\bar{v}_\delta) = 1$. Since $\phi_{K^*}(v)$ is strictly concave, its derivative is decreasing [25]. Therefore, on $[0, \bar{v}_\delta)$, it holds that

$$\phi'_{K^*}(v) \geq 1.$$

This implies that

$$\begin{aligned} \bar{v}_\delta &= \phi_{K^*}(\bar{v}_\delta) + \delta \\ &= \phi_{K^*}(0) + \int_0^{\bar{v}_\delta} \phi'_{K^*}(v)dx + \delta \\ &\geq \phi_{K^*}(0) + \delta + \bar{v}_\delta \\ &> \bar{v}_\delta. \end{aligned}$$

This is a contradiction, so it must be that $\phi'_{K^*}(\bar{v}_\delta) < 1$.

Since there is a unique positive fixed point $\bar{v}_\delta$ and $v - (\phi_{K^*}(v) + \delta) \to \infty$, it must hold that $\phi_{K^*}(x) + \delta \leq x$ iff $x \geq \bar{v}_\delta$. Since $\phi_{K^*}(\varepsilon) \leq \varepsilon$, it holds that $\bar{v}_0 \leq \varepsilon$.

Finally, for $\delta' \geq \delta$, it holds that

$$\begin{aligned} \bar{v}_\delta &= \phi_{K^*}(\bar{v}_\delta) + \delta \\ &= \phi_{K^*}(\bar{v}_\delta) + \delta' + \underbrace{(\delta - \delta')}_{<0} \\ &< \phi_{K^*}(\bar{v}_\delta) + \delta'. \end{aligned} \tag{37}$$

By the observation above, we then have $\bar{v}_\delta \leq \bar{v}_{\delta'}$. This monotonicity in turn implies that

$$\lim_{\delta \searrow 0} \bar{v}_\delta = \bar{v}_0.$$

□

As a simple consequence of the concavity of $\phi_{K^*}(v)$, we can study a fixed point iteration involving $\phi_K(v)$. Define the n-fold composition mapping

$$(\phi_K + \delta)^{(n)}(v) \triangleq ((\phi_K + \delta) \circ \cdots \circ (\phi_K + \delta))(v).$$

Lemma 7. *For any $v > 0$, it holds that*

$$\lim_{n\to\infty} (\phi_{K^*} + \delta)^{(n)}(v) = \bar{v}_\delta.$$

Proof. Following [30], for any fixed point $\bar{v}$, it holds that

$$|\phi_{K^*}(v) + \delta - \bar{v}_\delta| \leq \phi'_{K^*}(\bar{v})|v - \bar{v}_\delta|.$$

Therefore, applying the fixed point property repeatedly yields

$$|(\phi_{K^*} + \delta)^{(n)}(v) - \bar{v}_\delta| \leq (\phi'_{K^*}(\bar{v}))^n |v - \bar{v}_\delta|.$$

By Lemma 6, it holds that

$$\phi'_{K^*}(\bar{v}) < 1$$

and so the result follows. □

This implies that if we select K^* stochastic gradients at every time instant, and we start from any v, then it holds that

$$\phi^{(n)}_{K^*,\rho}(v) \to \bar{v}_0$$

with $\bar{v}_0 \leq \varepsilon$.

Now, we show that we control the mean criterion defined in eq. (5) when we estimate ρ. In Section 3.1, we pick a deterministic choice of $K_n = K^*$ and proceed with the analysis. Then it holds that

$$\begin{aligned}
\mathbb{E}[f_n(\boldsymbol{x}_n)] - f_n(\boldsymbol{x}_n^*) &\leq \mathbb{E}\left[b\left(\sqrt{\frac{2}{m}(f_{n-1}(\boldsymbol{x}_{n-1}) - f_{n-1}(\boldsymbol{x}^*_{n-1}))} + \rho, K_n\right)\right] \\
&= \mathbb{E}\left[\alpha(K_n)\left(\sqrt{\frac{2}{m}(f_{n-1}(\boldsymbol{x}_{n-1}) - f_{n-1}(\boldsymbol{x}^*_{n-1}))} + \rho\right)^2 + \beta(K_n)\right] \qquad (38) \\
&= \alpha(K_n)\mathbb{E}\left[\left(\sqrt{\frac{2}{m}(f_{n-1}(\boldsymbol{x}_{n-1}) - f_{n-1}(\boldsymbol{x}^*_{n-1}))} + \rho\right)^2\right] + \beta(K_n) \qquad (39) \\
&= \alpha(K^*)\mathbb{E}\left[\left(\sqrt{\frac{2}{m}(f_{n-1}(\boldsymbol{x}_{n-1}) - f_{n-1}(\boldsymbol{x}^*_{n-1}))} + \rho\right)^2\right] + \beta(K^*). \qquad (40)
\end{aligned}$$

We can bound

$$\mathbb{E}\left[\left(\sqrt{\frac{2}{m}(f_{n-1}(\boldsymbol{x}_{n-1}) - f_{n-1}(\boldsymbol{x}^*_{n-1}))} + \rho\right)^2\right]$$

using eq. (14) and recover eq. (35). However, in this paper, K_n and $\boldsymbol{x}_{n-1}$ are dependent random variables, so (39) does not hold in general. Instead, only (38) holds. To get around this issue, we need a more sophisticated analysis using the observation that $K_n \geq K^*$ for all n large enough. This property implies that K_n behaves like a constant for n large enough and the analysis in Section 3.1 nearly applies.

Proof of Theorem 5. We know that for all n large enough that we pick $K_n \geq K^*$ almost surely. This in turn implies that there exists a finite almost surely random variable $\tilde{N}$ such that

$$n \geq \tilde{N} \;\Rightarrow\; K_n \geq K^*.$$

Since $\tilde{N}$ is finite almost surely, we know that

$$\lim_{n\to\infty} \mathbb{P}\left\{\tilde{N} > n\right\} = 0.$$

By the compactness of $\mathscr{X}$, it follows that there is a constant $C > 0$ such that

$$\max_{\boldsymbol{x}\in\mathscr{X}} \phi_{K^*,\rho}\left(f_n(\boldsymbol{x}) - f_n(\boldsymbol{x}_n^*)\right) \le C \;\; \forall n \ge 1.$$

Then it follows that

$$\begin{aligned}
\mathbb{E}[f_n(\boldsymbol{x}_n)] - f_n(\boldsymbol{x}_n^*) &= \mathbb{E}\left[\phi_{K_n,\rho}\left(f_{n-1}(\boldsymbol{x}_{n-1}) - f_{n-1}(\boldsymbol{x}_{n-1}^*)\right)\right] \\
&= \mathbb{E}\left[\phi_{K_n,\rho}\left(f_{n-1}(\boldsymbol{x}_{n-1}) - f_{n-1}(\boldsymbol{x}_{n-1}^*)\right)\mathbb{1}_{\{n\ge\tilde{N}\}}\right] \\
&\quad + \mathbb{E}\left[\phi_{K_n,\rho}\left(f_{n-1}(\boldsymbol{x}_{n-1}) - f_{n-1}(\boldsymbol{x}_{n-1}^*)\right)\mathbb{1}_{\{n<\tilde{N}\}}\right] \\
&\le \mathbb{E}\left[\phi_{K^*,\rho}\left(f_{n-1}(\boldsymbol{x}_{n-1}) - f_{n-1}(\boldsymbol{x}_{n-1}^*)\right)\mathbb{1}_{\{n\ge\tilde{N}\}}\right] + C\mathbb{P}\left\{\tilde{N} > n\right\} \\
&\le \phi_{K^*,\rho}\left(\mathbb{E}[f_{n-1}(\boldsymbol{x}_{n-1})] - f_{n-1}(\boldsymbol{x}_{n-1}^*)\right) + C\mathbb{P}\left\{\tilde{N} > n\right\}.
\end{aligned}$$

To bound the mean criterion, we consider the recursion

$$\varepsilon_n = \phi_{K^*,\rho}\left(\varepsilon_{n-1}\right) + C\mathbb{P}\left\{\tilde{N} > n\right\} \qquad \forall n \ge \tilde{N} \tag{41}$$

which satisfies

$$\mathbb{E}[f_n(\boldsymbol{x}_n)] - f_n(\boldsymbol{x}_n^*) \le \varepsilon_n \qquad \forall n \ge \tilde{N}.$$

By assumption, we know that as $n \to \infty$

$$C\mathbb{P}\left\{\tilde{N} > n\right\} \to 0.$$

Fix $\delta > 0$. Then there exists a random variable $\tilde{N}_\delta \ge \tilde{N}$ such that

$$n \ge \tilde{N}_\delta \;\; \Rightarrow \;\; C\mathbb{P}\left\{\tilde{N} > n\right\} \le \delta.$$

Then we consider the recursion

$$\begin{aligned}
\tilde{\varepsilon}_n &= \phi_{K^*,\rho}\left(\tilde{\varepsilon}_{n-1}\right) + \delta \\
\tilde{\varepsilon}_{\tilde{N}_\delta} &= \varepsilon_{\tilde{N}_\delta}
\end{aligned} \qquad \forall n \ge \tilde{N}_\delta. \tag{42}$$

By construction, we have $\varepsilon_n \le \tilde{\varepsilon}_n$ for all $n \ge \tilde{N}_\delta$. As a consequence of Lemma 6 and Lemma 7, we have

$$\begin{aligned}
\limsup_{n\to\infty}\left(\mathbb{E}[f_n(\boldsymbol{x}_n)] - f_n(\boldsymbol{x}_n^*)\right) &\le \limsup_{n\to\infty} \varepsilon_n \\
&\le \limsup_{n\to\infty} \tilde{\varepsilon}_n \\
&\le \overline{v}_\delta.
\end{aligned}$$

Since $\delta > 0$ was arbitrary and $\overline{v}_\delta \searrow \overline{v}_0$ as $\delta \searrow 0$ from Lemma 6, it follows that

$$\limsup_{n\to\infty}\left(\mathbb{E}[f_n(\boldsymbol{x}_n)] - f_n(\boldsymbol{x}_n^*)\right) \le \overline{v}_0 \le \varepsilon.$$

□

10 Examples of $b(d_0, K)$ Bound for SGD

We examine bounds $b(d_0, K)$ satisfying assumption A.4 for SGD Equation (10). We form a convex combination of the iterates to yield a final approximate minimizer

$$\overline{\boldsymbol{x}}(K) \;\; = \;\; \sum_{k=0}^{K} \lambda(k)\boldsymbol{x}(k).$$

Note that this includes the case where $\overline{\boldsymbol{x}}(K) = \boldsymbol{x}(K)$ by selecting $\lambda(K) = 1$ and $\lambda(0) = \cdots = \lambda(K-1) = 0$.

Define

$$d(k) \triangleq \|\boldsymbol{x}(k) - \boldsymbol{x}^*\|_2. \tag{43}$$

First, we bound $\mathbb{E}[d(k)]$ in Lemma 8, which follows the classic Lyapunov function analysis of SGD [31].

Lemma 8. *It holds that*

$$\mathbb{E}[d(k)] \leq \prod_{k=1}^{K}(1-2m\mu(k)+B\mu^2(k))d^2(0)+A\sum_{k=1}^{K}\prod_{i=k+1}^{K}(1-2m\mu(i)+B\mu^2(i))\mu^2(k).$$

Proof. See [31]. □

It is possible to further upper bound the bound in Lemma 8 to yield a closed form given in [32]; however, the bound in Lemma 8 is generally tighter. Next, we apply Lemma 8 along with a Lipschitz gradient assumption on $f(\boldsymbol{x})$ to produce a simple $b(d_0,K)$ bound.

Lemma 9. *With arbitrary step sizes, assuming that $f(\boldsymbol{x})$ has Lipschitz continuous gradients with modulus M, and $\lambda(K)=1$, it holds that*

$$\mathbb{E}[f(\bar{\boldsymbol{x}}(K))]-f(\boldsymbol{x}^*)\leq\frac{1}{2}M\mathbb{E}[d^2(K)]$$

and therefore, it holds that

$$b(d_0,K)=\frac{1}{2}M\left(\prod_{k=1}^{K}(1-2m\mu(k)+B\mu^2(k))d_0^2+A\sum_{k=1}^{K}\prod_{i=k+1}^{K}(1-2m\mu(i)+B\mu^2(i))\mu^2(k)\right)$$

satisfies the requirements of assumption A.4.

Proof. Using the descent lemma from [33], it holds that
$\mathbb{E}[f(\boldsymbol{x})]-f(\boldsymbol{x}^*)\leq\frac{1}{2}M\mathbb{E}[d(K)]$. Plugging in the bound from Lemma 8 yields the bound $b(d_0,K)$. □

Next, we consider an extension of the averaging scheme derived with $B=0$ in [34] to the case with $B>0$ using the bounds in Lemma 8. This averaging scheme puts weight

$$\lambda(k)=\frac{\frac{1}{\mu(k)}}{\sum_{j=1}^{K}\frac{1}{\mu(j)}}$$

on the iterate $\boldsymbol{x}(k)$ with step size $\mu(k)=\mathcal{O}(k^{-1})$. Therefore, this averaging puts increasing weight on later iterates.

Lemma 10. *With the choice of step sizes given by*

$$\mu(k)=\frac{1}{m(k+1)}\quad\forall k\geq 1$$

it holds that

$$b(d_0,K)=\frac{(1+B)d_0^2+B\sum_{k=1}^{K}\gamma(k)+(K+1)A}{m(K+1)(K+4)}$$

satisfies assumption A.4 where $\mathbb{E}[d(k)]\leq\gamma(k)$.

Proof. This proof is a straightforward extension of the proof in [34]. We have using standard analysis of SGD (see [31] for example)

$$\mathbb{E}[d^2(k)]\leq(1-2m\mu(k)+B\mu^2(k))\mathbb{E}[d^2(k-1)]-2\mu(k)(\mathbb{E}[f(\boldsymbol{x}(k-1))]-f(\boldsymbol{x}^*))+A\mu^2(k).$$

Then dividing by $\mu^2(k)$, we have

$$\frac{1}{\mu^2(k)}\mathbb{E}[d^2(k)]\leq\left(\frac{1-2m\mu(k)}{\mu^2(k)}+B\right)\mathbb{E}[d^2(k-1)]-\frac{2}{\mu(k)}(\mathbb{E}[f(\boldsymbol{x}(k-1))]-f(\boldsymbol{x}^*))+A.$$

It holds that

$$\frac{1-2m\mu(k)}{\mu^2(k)}\leq\frac{1}{\mu^2(k-1)}.$$

This implies that

$$\frac{1}{\mu^2(k)}\mathbb{E}[d(k)] - \frac{1}{\mu^2(k-1)}\mathbb{E}[d^2(k-1)] \leq B\mathbb{E}[d^2(k-1)] - \frac{2}{\mu(k)}(\mathbb{E}[f(\boldsymbol{x}(k-1))] - f(\boldsymbol{x}^*)) + A.$$

Summing from $k=1$ to $K+1$ and rearranging yields

$$\sum_{k=0}^{K} \frac{1}{\mu(k+1)} (\mathbb{E}[f(\boldsymbol{x}(k))] - f(\boldsymbol{x}^*)) \leq \frac{1}{2}(1+B)d_0^2 + \frac{1}{2}B\sum_{k=1}^{K}\mathbb{E}[d(k)] + \frac{1}{2}(K+1)A.$$

With the weights

$$\lambda(k) = \frac{\frac{1}{\mu(k+1)}}{\sum_{\tau=0}^{K}\frac{1}{\mu(j+1)}}$$

we have

$$\mathbb{E}[f(\bar{\boldsymbol{x}}(K))] - f(\boldsymbol{x}^*) \leq \frac{\frac{1}{2}(1+B)d(0) + \frac{1}{2}B\sum_{k=1}^{K}\mathbb{E}[d^2(k)] + \frac{1}{2}(K+1)A}{\sum_{\tau=0}^{K}\frac{1}{\mu(\tau)}}.$$

Then it holds that

$$\sum_{\tau=0}^{K}\frac{1}{\mu(\tau+1)} = \sum_{\tau=0}^{K} m(\tau+2) = \frac{1}{2}m(K+1)(K+4)$$

so

$$\begin{aligned}\mathbb{E}[f(\bar{\boldsymbol{x}}(K))] - f(\boldsymbol{x}^*) &\leq \frac{(1+B)d(0) + B\sum_{k=1}^{K}\mathbb{E}[d(k)] + (K+1)A}{m(K+1)(K+4)} \\ &\leq \frac{(1+B)d_0^2 + B\sum_{k=1}^{K}\gamma(k) + (K+1)A}{m(K+1)(K+4)}.\end{aligned}$$

□

To get the required $\gamma(k)$ bounds, we use Lemma 8. For the choice of step sizes in Lemma 10 from Lemma 8, it holds that $\mathbb{E}[d(k)] = \mathscr{O}\left(\frac{1}{k}\right)$. Since

$$\sum_{k=1}^{K}\frac{1}{k} = \mathscr{O}(\log K)$$

it holds that

$$\mathbb{E}[f(\bar{\boldsymbol{x}}(K))] - f(\boldsymbol{x}^*) = \mathscr{O}\left(\frac{d_0^2}{K^2} + \frac{\log(K)}{K^2} + \frac{1}{K}\right).$$

The $\mathscr{O}(\frac{1}{K})$ rate is minimax optimal for stochastic minimization of a strongly convex function [35].

Next, we look at a special case of averaging from [32] for stochastic gradients such that

$$\mathbb{E}\|\boldsymbol{g}(\boldsymbol{x},\boldsymbol{z}) - \boldsymbol{g}(\tilde{\boldsymbol{x}},\boldsymbol{z}) - \boldsymbol{g}^{(2)}(\boldsymbol{x},\boldsymbol{z})(\boldsymbol{x}-\tilde{\boldsymbol{x}})\|_2^2 = 0 \quad \forall \boldsymbol{x},\tilde{\boldsymbol{x}} \in \mathscr{X}$$

where $\boldsymbol{g}^{(2)}(\boldsymbol{x},\boldsymbol{z})$ is an unbiased stochastic second derivative with respect to $\boldsymbol{x}$. Quadratic objectives satisfy this condition.

Lemma 11. *Assuming that*

1. $\mathbb{E}\|\boldsymbol{g}(\boldsymbol{x},\boldsymbol{z}) - \boldsymbol{g}(\tilde{\boldsymbol{x}},\boldsymbol{z}) - \boldsymbol{g}^{(2)}(\boldsymbol{x},\boldsymbol{z})(\boldsymbol{x}-\tilde{\boldsymbol{x}})\|_2^2 = 0 \quad \forall \boldsymbol{x},\tilde{\boldsymbol{x}} \in \mathscr{X}$
2. $\mu(k) = Ck^{-\alpha}$ *with* $\alpha \geq 1/2$
3. $\lambda(0) = 0$ *and* $\lambda(k) = 1/K$ *for* $1 \leq k \leq K$
4. $\mathbb{E}[d^2(k)] \leq \gamma(k)$

it holds that

$$\left(\mathbb{E}[\bar{d}(K)]\right)^{1/2} \leq \frac{1}{m^{1/2}} \sum_{k=1}^{K-1} \left| \frac{1}{\mu(k+1)} - \frac{1}{\mu(k)} \right| (\gamma(k))^{1/2} + \frac{1}{m^{1/2}\mu(1)} \left(d_0^2\right)^{1/2} + \frac{1}{m^{1/2}\mu(K)} (\gamma(K))^{1/2} + \sqrt{\frac{A}{mK}} + \sqrt{\frac{2B}{mK^2} \sum_{k=1}^{K} \mathbb{E}[d^2(k-1)]}$$

with $\bar{d}(K) = \|\bar{\boldsymbol{x}}(K) - \boldsymbol{x}^*\|_2^2$. *If* f *has Lipschitz continuous gradients with modulus* M*, then it holds that*

$$b(d_0, K) = \frac{M}{2} \left(\frac{1}{m^{1/2}} \sum_{k=1}^{K-1} \left| \frac{1}{\mu(k+1)} - \frac{1}{\mu(k)} \right| (\gamma(k))^{1/2} + \frac{1}{m^{1/2}\mu(1)} (d(0))^{1/2} + \frac{1}{m^{1/2}\mu(K)} (\gamma(K))^{1/2} + \sqrt{\frac{A}{mK}} + \sqrt{\frac{2B}{mK^2} \sum_{k=1}^{K} \mathbb{E}[d(k-1)]} \right)^2$$

satisfies assumption A.4.

Proof. See [32] for the proof □

This decays at rate $\mathcal{O}\left(\frac{1}{K}\right)$ as long as $\mu(k) = Ck^{-\alpha}$ with $\frac{1}{2} \leq \alpha \leq 1$. To get the bounds $\gamma(k)$, we can again apply Lemma 8.

11 IHP Tracking Error Analysis

For the IHP criterion, we assume that assumptions A.1-A.6 hold. We seek an upper bound $r(t,K)$ such that

$$\mathbb{P}\{f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t\} \leq r(t,K) \quad \forall n \geq 1. \tag{44}$$

Using the mean criterion bounds of the previous section, we know that for all n

$$\mathbb{E}[f_n(\boldsymbol{x}_n)] - f_n(\boldsymbol{x}_n^*) \leq \varepsilon.$$

Then by Markov's inequality, it holds that

$$\mathbb{P}\{f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t\} \leq \frac{\varepsilon}{t}. \tag{45}$$

Although this bound always holds, we look at a way to tighten this bound. As before, we proceed by induction. As a base case, we can set

$$\mathbb{P}\{f_1(\boldsymbol{x}_1) - f_1(\boldsymbol{x}_1^*) > t\} \leq \frac{\varepsilon}{t}.$$

Now, suppose that

$$\mathbb{P}\left\{f_{n-1}(\boldsymbol{x}_{n-1}) - f_{n-1}(\boldsymbol{x}_{n-1}^*) > t\right\} \leq r_{n-1}(t)$$

and we want to construct a bound $r_n(t)$ on $\mathbb{P}\{f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t\}$. We proceed by conditioning on $\{d_n(0) \leq \delta\}$ and $\{d_n(0) > \delta\}$ with $d_n(0)$ defined in eq. (12) using the law of total probability:

$$\begin{aligned} \mathbb{P}\{f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t\} \quad = \quad & \mathbb{P}\{f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t \mid d_n(0) \leq \delta\} \mathbb{P}\{d_n(0) \leq \delta\} \\ & + \mathbb{P}\{f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t \mid d_n(0) > \delta\} \mathbb{P}\{d_n(0) > \delta\}. \end{aligned}$$

For the first term, it holds that

$$\mathbb{P}\{f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t \mid d_n(0) \leq \delta\} \leq \frac{1}{t} b(\delta, K) \triangleq \psi(t, \delta) \tag{46}$$

and

$$\mathbb{P}\{d_n(0) \leq \delta\} \leq 1.$$

For the second term, it holds that

$$\mathbb{P}\{f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t \mid d_n(0) > \delta\} \leq \psi(t, \mathrm{diam}(\mathscr{X}))$$

and

$$\begin{aligned}
\mathbb{P}\{d_n(0) > \delta\} &= \mathbb{P}\left\{\|\boldsymbol{x}_{n-1} - \boldsymbol{x}_n^*\|_2^2 > \delta\right\} \\
&= \mathbb{P}\left\{\|\boldsymbol{x}_{n-1} - \boldsymbol{x}_n^*\|_2 > \sqrt{\delta}\right\} \\
&\leq \mathbb{P}\left\{\|\boldsymbol{x}_{n-1} - \boldsymbol{x}_{n-1}^*\|_2 + \rho > \sqrt{\delta}\right\} \\
&\leq \mathbb{P}\left\{\|\boldsymbol{x}_{n-1} - \boldsymbol{x}_{n-1}^*\|_2 > \left(\sqrt{\delta} - \rho\right)_+\right\} \\
&\leq \mathbb{P}\left\{\sqrt{f_{n-1}(\boldsymbol{x}_{n-1}) - f_{n-1}(\boldsymbol{x}_{n-1}^*)} > \sqrt{\frac{2}{m}}\left(\sqrt{\delta} - \rho\right)_+\right\} \\
&\leq r_{n-1}\left(\frac{2}{m}\left(\sqrt{\delta} - \rho\right)_+^2\right)
\end{aligned}$$

where $(x)_+ = \max\{x, 0\}$. Combining these bounds yields an overall bound

$$\mathbb{P}\{f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t\} \leq \psi(t, \delta) + \psi(t, \mathrm{diam}(\mathscr{X}))r_{n-1}\left(\frac{2}{m}\left(\sqrt{\delta} - \rho\right)_+^2\right).$$

We can optimize this bound over δ to yield the bound

$$\begin{aligned}
&\mathbb{P}\{f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t\} \\
&\quad \leq \inf_{0 < \delta \leq \mathrm{diam}(\mathscr{X})}\left\{\psi(t, \delta) + \psi(t, \mathrm{diam}(\mathscr{X}))r_{n-1}\left(\frac{2}{m}\left(\sqrt{\delta} - \rho\right)_+^2\right)\right\}. \qquad (47)
\end{aligned}$$

The quantity $\psi(t, \delta)$ defined in eq. (46) can be replaced by any bound that also satisfies the inequality in eq. (46). Therefore, we can set

$$r_n(t) = \inf_{0 < \delta \leq \mathrm{diam}(\mathscr{X})}\left\{\psi(t, \delta) + \psi(t, \mathrm{diam}(\mathscr{X}))r\left(\frac{2}{m}\left(\sqrt{\delta} - \rho\right)_+^2\right)\right\} \qquad (48)$$

resulting in $\mathbb{P}\{f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t\} \leq r_n(t)$. In Section 11.1, we develop an approach to compute the IHP bound. With this approach, we can set K_n to achieve a desired IHP criterion ε and r.

11.1 Bound at a Finite Number of Points

The bound of the preceding section is exact but difficult to compute. In this section, we introduce a computationally simpler bound. Computing the entire sequence of functions $r_n(t)$ is generally difficult, so we look at bounding

$$\mathbb{P}\{f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t\}$$

at a finite number of points $t^{(1)}, \ldots, t^{(N)}$ ordered in increasing order. We want to compute bound $r_n(1), \ldots, r_n(N)$ such that

$$\mathbb{P}\left\{f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t^{(i)}\right\} \leq r_n(i) \quad i = 1, \ldots, N.$$

We define an initial bound

$$r_1(i) = \frac{\varepsilon}{t^{(i)}}.$$

Suppose that

$$\mathbb{P}\left\{ f_{n-1}(\boldsymbol{x}_{n-1}) - f_{n-1}(\boldsymbol{x}^*_{n-1}) > t^{(i)} \right\} \leq r_{n-1}(i).$$

Then as in eq. (48), it follows that

$$\mathbb{P}\left\{ f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}^*_n) > t^{(i)} \right\} \leq \inf_{0<\delta\leq \mathrm{diam}(\mathscr{X})} \left\{ \psi(t^{(i)},\delta) + \psi(t^{(i)},\mathrm{diam}(\mathscr{X}))\mathbb{P}\{d_n(0)>\delta\} \right\}. \tag{49}$$

The key then is to bound $\mathbb{P}\{d_n(0)>\delta\}$ in terms of $\{r_{n-1}(i)\}_{i=1}^N$. Define the function

$$t(\delta) = \max\{t^{(i)} \mid t^{(i)} \leq \delta\}$$

to be point $t^{(i)}$ closest to δ but not greater. Provided that $\frac{2}{m}\left(\sqrt{\delta}-\rho\right)_+^2 \geq t^{(1)}$ and $t^{(i)} = t\left(\frac{2}{m}\left(\sqrt{\delta}-\rho\right)_+^2\right)$ it holds that

$$\begin{aligned}
\mathbb{P}\{d_n(0)>\delta\} &\leq \mathbb{P}\left\{ f_{n-1}(\boldsymbol{x}_{n-1}) - f_{n-1}(\boldsymbol{x}^*_{n-1}) > \frac{2}{m}\left(\sqrt{\delta}-\rho\right)_+^2 \right\} \\
&\leq \mathbb{P}\left\{ f_{n-1}(\boldsymbol{x}_{n-1}) - f_{n-1}(\boldsymbol{x}^*_{n-1}) > t^{(i)} \right\} \\
&\leq r_{n-1}(i).
\end{aligned} \tag{50}$$

Otherwise, if $\frac{2}{m}\left(\sqrt{\delta}-\rho\right)_+^2 < t^{(1)}$, then

$$\mathbb{P}\{d_n(0)>\delta\} \leq \frac{\varepsilon}{\frac{2}{m}\left(\sqrt{\delta}-\rho\right)_+}.$$

Define the overall bound for the term $\mathbb{P}\{d_n(0)>\delta\}$ as follows:

$$\phi_n(\delta) \triangleq \begin{cases} r_{n-1}\left(t\left(\frac{2}{m}\left(\sqrt{\delta}-\rho\right)_+^2\right)\right), & t^{(1)} \leq \frac{2}{m}\left(\sqrt{\delta}-\rho\right)_+^2 \\ \frac{\varepsilon}{\frac{2}{m}\left(\sqrt{\delta}-\rho\right)_+^2}, & \text{else} \end{cases}. \tag{51}$$

Then we can set

$$r_n^{(i)} = \inf_{\delta>0}\left\{ \psi(t^{(i)},\delta) + \psi(t^{(i)},\mathrm{diam}(\mathscr{X}))\phi_n(\delta) \right\}.$$

This algorithm is summarized in Algorithm algorithm 1. In practice, once the bound $\psi(t^{(i)},\mathrm{diam}(\mathscr{X}))$ is less than one, then the gains are significant. Figure fig. 5 plots a comparison of the bound produced by Algorithm algorithm 1 against the Markov inequality bound from eq. (45) applied to the problem in Section 6 with $K_n = 300$.

Algorithm 1 Calculate IHP bounds

Require: Points $t^{(1)},\ldots,t^{(N)}$

1. Set
$$r_1^{(i)} = \frac{\varepsilon}{t^{(i)}} \quad i = 1,\ldots,N.$$

2. Compute
$$r_n^{(i)} = \min_{0<\delta\leq \mathrm{diam}(\mathscr{X})}\left\{ \psi(t^{(i)},\delta) + \psi(t^{(i)},\mathrm{diam}(\mathscr{X}))\phi_n(\delta) \right\}$$
for $i = 1,\ldots,N$.

3. $n \leftarrow n+1$ and go back to step 2.

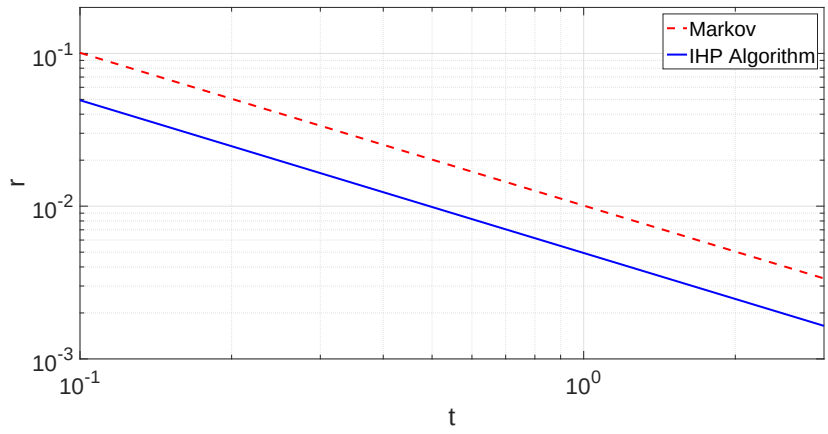

Figure 5: IHP Algorithm Plot

Either the Markov bound of eq. (45) or Algorithm algorithm 1 will produce valid upper bounds of the form

$$\mathbb{P}\left\{ f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t^{(i)} \right\} \leq r_n(i) \quad i = 1, \ldots, N.$$

Suppose that the set $\{t^{(1)}, \ldots, t^{(N)}\}$ contains t at index i^*. These bounds can in turn be used to select K_n to achieve a target (t, r) pair by selecting the smallest K_n such that

$$r_n(i^*) \leq r.$$

11.2 In High Probability Bounds With Change in Minimizer Unknown

We can adopt the same approach as with ρ known by substituting $\hat{\rho}_{n-1} + t_{n-1}$ in place of ρ. As soon as $\rho_{n-1} + t_{n-1} \geq \rho$, algorithm 1 will produce upper bounds of the form

$$\mathbb{P}\left\{ f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t^{(i)} \right\} \leq r_n(i) \quad i = 1, \ldots, N.$$

Suppose that the set $\{t^{(1)}, \ldots, t^{(N)}\}$ contains t at index i^*. These bounds can in turn be used to select K_n to achieve a target (t, r) pair by selecting the smallest K_n such that

$$r_n(i^*) \leq r.$$

12 More Proofs for Estimates of Change in Minimizers

We can also prove the result for the inequality constraint of eq. (2) using similar techniques in Theorem 3.

Proof of Theorem 3. Define $\bar{\rho}_i^{(2)}$ and $\bar{\rho}_i^{(3)}$ analogous to the equality case proof and the pair $\hat{\rho}_i^{(2)}$ and $\hat{\rho}_i^{(3)}$ of the form

in eq. (26). First, we have by assumptions B.4-B.5

$$\begin{aligned}
|\hat{\rho}_n - \hat{\rho}_n^{(3)}| &\leq \frac{1}{n-W}\sum_{i=W+1}^{n} |\overline{\rho}^{(i)} - \overline{\rho}_3^{(i)}| \\
&\leq \frac{1}{n-W}\sum_{i=W+1}^{n}\sum_{j=i-W+1}^{i} a_j |\tilde{\rho}_j - \tilde{\rho}_j^{(3)}| \\
&\leq \frac{1}{n-W}\sum_{i=W+1}^{n}\sum_{j=i-W+1}^{i} a_j \left(|\tilde{\rho}_j - \tilde{\rho}_j^{(2)}| + |\tilde{\rho}_j^{(2)} - \tilde{\rho}_j^{(3)}|\right) \\
&\leq \frac{\sum_{j=1}^{W} a_j}{n-W}\sum_{i=2}^{n}\left(|\tilde{\rho}_i - \tilde{\rho}_i^{(2)}| + |\tilde{\rho}_i^{(2)} - \tilde{\rho}_i^{(3)}|\right) \\
&\leq \left(\frac{n-1}{n-W}\sum_{j=1}^{W} a_j\right)\frac{1}{n-1}\sum_{i=2}^{n}\left(|\tilde{\rho}_i - \tilde{\rho}_i^{(2)}| + |\tilde{\rho}_i^{(2)} - \tilde{\rho}_i^{(3)}|\right).
\end{aligned}$$

In comparison, we looked at controlling

$$\frac{1}{n-1}\sum_{i=2}^{n}\left(|\tilde{\rho}_i - \tilde{\rho}_i^{(2)}| + |\tilde{\rho}_i^{(2)} - \tilde{\rho}_i^{(3)}|\right)$$

in the proof of Theorem 1. The quantity of interest here is the same scaled by

$$\frac{n-1}{n-W}\sum_{j=1}^{W} a_j.$$

By construction, we always have $\tilde{\rho}_i^{(3)} \geq \rho_i$. Therefore, by assumptions B.4-B.5, it follows that

$$\mathbb{E}\left[h(\tilde{\rho}_i^{(3)}, \tilde{\rho}_{i-1}^{(3)}, \ldots, \tilde{\rho}_{i-W+1}^{(3)}) \mid \mathscr{F}_{i-1}\right] \geq \rho_i$$

and $\mathbb{E}[\hat{\rho}_n^{(3)} \mid \mathscr{F}_{i-1}] \geq \rho$. Therefore, by applying Lemma 3 and the Borel-Cantelli lemma, it follows for all n large enough

$$\hat{\rho}_n^{(3)} + t_n \geq \rho.$$

This observation combined with a nearly identical proof to equality case shows that for all n large enough and appropriate $\{t_n\}$

$$\hat{\rho}_n + \left(\frac{n-1}{n-W}\sum_{j=1}^{W} a_j\right) D_n + t_n \geq \rho$$

almost surely. □

12.1 L_2 Norm Condition

Now, we look at analyzing $\hat{\rho}_n$ from eq. (24) under the L_2 condition. First, we consider the condition in eq. (17). Define the averaged estimate

$$\left(\hat{\rho}_n^{(3)}\right)^2 \triangleq \frac{1}{n-1}\sum_{i=2}^{n}\left(\tilde{\rho}_i^{(3)}\right)^2$$

and analogously $\left(\hat{\rho}_n^{(2)}\right)^2$. The following lemma shows that $\hat{\rho}_n^{(3)}$ upper bounds ρ eventually.

Lemma 12. *For all sequences $\{t_n\}$ such that*

$$\sum_{n=2}^{\infty} \exp\left\{-\frac{2(n-1)t_n^2}{diam^2(\mathscr{X})}\right\} < +\infty$$

it holds that for all n large enough

$$\sqrt{\left(\hat{\rho}_n^{(3)}\right)^2 + t_n} \geq \rho$$

almost surely.

Proof. First, for all i it holds that

$$\tilde{\rho}_i^{(3)} \geq \|\boldsymbol{x}_i^* - \boldsymbol{x}_{i-1}^*\|.$$

This in turn implies that

$$\mathbb{E}\left[\left(\tilde{\rho}_i^{(3)}\right)^2 \,\middle|\, \mathscr{F}_{i-1}\right] \geq \mathbb{E}\left[\|\boldsymbol{x}_i^* - \boldsymbol{x}_{i-1}^*\|^2 \mid \mathscr{F}_{i-1}\right] = \rho^2.$$

Second, it holds that $0 \leq \left(\tilde{\rho}_i^{(3)}\right)^2 \leq \text{diam}^2(\mathscr{X})$. Applying Lemma 4 and Lemma 3 yields

$$\begin{aligned}
\mathbb{P}\left\{\left(\hat{\rho}_n^{(3)}\right)^2 < \rho^2 - t_n\right\} &\leq \mathbb{P}\left\{\left(\hat{\rho}_n^{(3)}\right)^2 < \frac{1}{n-1}\sum_{i=2}^{n}\mathbb{E}\left[\left(\tilde{\rho}_i^{(3)}\right)^2 \,\middle|\, \mathscr{F}_i\right] - t_n\right\} \\
&\leq \exp\left\{-\frac{2(n-1)t_n^2}{\text{diam}^4(\mathscr{X})}\right\}.
\end{aligned}$$

By the Borel-Cantelli lemma, this in turn implies that for n sufficiently large

$$\sqrt{\left(\hat{\rho}_n^{(3)}\right)^2 + t_n} \geq \rho.$$

□

We can now follow the proof technique of Theorem 1 and Lemma 12 to prove Theorem 2.

Proof of Theorem 2. This is a straightforward extension of the proof of Theorem 1 using the observation that we have

$$\begin{aligned}
|(\hat{\rho}_n)^2 - (\hat{\rho}_n^{(3)})^2| &\leq |(\hat{\rho}_n)^2 - (\hat{\rho}_n^{(2)})^2| + |(\hat{\rho}_n^{(2)})^2 - (\hat{\rho}_n^{(3)})^2| \\
&\leq |\hat{\rho}_n + \hat{\rho}_n^{(2)}||\hat{\rho}_n - \hat{\rho}_n^{(2)}| + |\hat{\rho}_n^{(2)} + \hat{\rho}_n^{(3)}||\hat{\rho}_n^{(2)} - \hat{\rho}_n^{(3)}| \\
&\leq 2\text{diam}(\mathscr{X})\left(|\hat{\rho}_n - \hat{\rho}_n^{(2)}| + |\hat{\rho}_n^{(2)} - \hat{\rho}_n^{(3)}|\right).
\end{aligned}$$

We can now follow the proof technique of Theorem 1. □

Proof of Theorem 4. This is a straightforward extension of the proof of Theorem 3 along the lines of Theorem 2. □

13 Parameter Estimation

We may need to estimate parameters of the functions $\{f_n\}$ such as the strong convexity parameter m to compute the bound $b(d_0, K)$ from assumption A.4. In this section, we assume that the bound $b(d_0, K, \psi)$ is parameterized by ψ, which depends on properties of the functions $f_n(\boldsymbol{x})$. In most cases, we have parameters

$$\psi = \begin{bmatrix} 1/m & A & B \end{bmatrix}^\top$$

where m is the parameter of strong convexity, and the pair (A, B) controls gradient growth as in assumption A.5, i.e.,

$$\mathbb{E}\|\boldsymbol{g}(\boldsymbol{x}, \boldsymbol{z})\|_2^2 \leq A + B\|\boldsymbol{x} - \boldsymbol{x}^*\|_2^2.$$

We parameterize using $1/m$, since smaller m increase the bound $b(d_0, K, \psi)$. Our goal is to produce an estimate $\hat{\psi}$ such that $\hat{\psi} \geq \psi^*$ with ψ^* the true parameters. We present several general methods for estimating these parameters..

Similar to estimating ρ, we produce one time instant estimates $\tilde{m}_i$, $\tilde{A}_i$, and $\tilde{B}_i$ at time i and combine them by averaging to yield

1. $\hat{m}_n = \frac{1}{n}\sum_{i=1}^n \tilde{m}_i$
2. $\hat{A}_n = \frac{1}{n}\sum_{i=1}^n \tilde{A}_i$
3. $\hat{B}_n = \frac{1}{n}\sum_{i=1}^n \tilde{B}_i$.

We make the following assumptions for our analysis:

D.1 The parameters $\psi \in \mathscr{P}$ with $\mathscr{P}$ compact and there exists a true set of parameters ψ^*.

D.2 The bound $b(d_0, K, \tilde{\psi})$ is non-decreasing in ψ, i.e.,

$$\psi \leq \tilde{\psi} \;\Rightarrow\; b(d_0, K, \psi) \leq b(d_0, K, \tilde{\psi}).$$

D.3 $\nabla f_n(\boldsymbol{x}_n)$ has Lipschitz continuous gradients with modulus M.

D.4 $f_n(\boldsymbol{x})$ is twice differentiable and there exist stochastic second derivatives with respect to $\boldsymbol{x}$, $\boldsymbol{g}^{(2)}(\boldsymbol{x}, \boldsymbol{z})$, such that

$$\mathbb{E}_{\boldsymbol{z}_n \sim p_n}\left[\boldsymbol{g}_n^{(2)}(\boldsymbol{x}, \boldsymbol{z}_n) \mid \boldsymbol{x}\right] = \nabla^2_{\boldsymbol{x}\boldsymbol{x}} f_n(\boldsymbol{x}).$$

D.5 The space $\mathscr{X}$ is compact and there exists a constant G such that $\|\boldsymbol{g}_n(\boldsymbol{x}, \boldsymbol{z})\|_2 \leq G \quad \forall \boldsymbol{x}\, \forall \boldsymbol{z}\, \forall n$.

D.6 We have access to stochastic functions $\hat{f}_n(\boldsymbol{x}, \boldsymbol{z})$ such that

$$\mathbb{E}[\hat{f}_n(\boldsymbol{x}, \boldsymbol{z}) \mid \boldsymbol{x}] = f_n(\boldsymbol{x}).$$

13.1 Estimating the Strong Convexity Parameter

We seek one step estimates $\tilde{m}_i$ of the parameter of the strong convexity such that

$$\mathbb{E}[\tilde{m}_i \mid \mathscr{F}_{i-1}] \leq m.$$

For any two points $\boldsymbol{x}$ and $\tilde{\boldsymbol{x}}$, by strong convexity we have

$$f_i(\tilde{\boldsymbol{x}}) \geq f_i(\boldsymbol{x}) + \langle \nabla f_i(\boldsymbol{x}), \tilde{\boldsymbol{x}} - \boldsymbol{x}\rangle + \frac{1}{2} m\|\tilde{\boldsymbol{x}} - \boldsymbol{x}\|_2^2 \quad \forall \boldsymbol{x}, \tilde{\boldsymbol{x}} \in \mathscr{X}$$

which implies that

$$m \leq \frac{f_i(\tilde{\boldsymbol{x}}) - f_i(\boldsymbol{x}) - \langle \nabla f_i(\boldsymbol{x}), \tilde{\boldsymbol{x}} - \boldsymbol{x}\rangle}{\frac{1}{2}\|\tilde{\boldsymbol{x}} - \boldsymbol{x}\|_2^2}.$$

We suppose that for all n

$$m = \min_{\boldsymbol{x}, \tilde{\boldsymbol{x}} \in \mathscr{X}} \frac{f_i(\tilde{\boldsymbol{x}}) - f_i(\boldsymbol{x}) - \langle \nabla f_i(\boldsymbol{x}), \tilde{\boldsymbol{x}} - \boldsymbol{x}\rangle}{\frac{1}{2}\|\tilde{\boldsymbol{x}} - \boldsymbol{x}\|_2^2}.$$

This is not restrictive since any $m > 0$ that satisfies

$$m \leq \min_{\boldsymbol{x}, \tilde{\boldsymbol{x}} \in \mathscr{X}} \frac{f_i(\tilde{\boldsymbol{x}}) - f_i(\boldsymbol{x}) - \langle \nabla f_i(\boldsymbol{x}), \tilde{\boldsymbol{x}} - \boldsymbol{x}\rangle}{\frac{1}{2}\|\tilde{\boldsymbol{x}} - \boldsymbol{x}\|_2^2}$$

can be taken as a parameter of strong convexity for the class of functions $f_i(\boldsymbol{x})$. We estimate this quantity for fixed $\boldsymbol{x}$ and $\tilde{\boldsymbol{x}}$ using the plug in approximation in eq. (52).

$$r(\boldsymbol{x},\tilde{\boldsymbol{x}}) \triangleq \frac{\frac{1}{K_i}\sum_{k=1}^{K_i}\hat{f}_i(\tilde{\boldsymbol{x}},\boldsymbol{z}_i(k)) - \frac{1}{K_i}\sum_{k=1}^{K_i}\hat{f}_i(\boldsymbol{x},\boldsymbol{z}_i(k)) - \left\langle \frac{1}{K_i}\sum_{k=1}^{K_i}\boldsymbol{g}_i(\boldsymbol{x},\boldsymbol{z}_i(k)),\tilde{\boldsymbol{x}}-\boldsymbol{x}\right\rangle}{\frac{1}{2}\|\tilde{\boldsymbol{x}}-\boldsymbol{x}\|_2^2}. \tag{52}$$

Then consider the following estimate of m:

$$\tilde{m}_i \triangleq \min_{\boldsymbol{x},\tilde{\boldsymbol{x}}\in\mathscr{X}} r(\boldsymbol{x},\tilde{\boldsymbol{x}}). \tag{53}$$

This estimate satisfies

$$\begin{aligned}
\mathbb{E}[\tilde{m}_i \mid \mathscr{F}_{i-1}] &= \mathbb{E}\left[\min_{\boldsymbol{x},\tilde{\boldsymbol{x}}\in\mathscr{X}} r(\boldsymbol{x},\tilde{\boldsymbol{x}}) \mid \mathscr{F}_{i-1}\right] \\
&\leq \min_{\boldsymbol{x},\tilde{\boldsymbol{x}}\in\mathscr{X}} \mathbb{E}[r(\boldsymbol{x},\tilde{\boldsymbol{x}})] \\
&= \min_{\boldsymbol{x},\tilde{\boldsymbol{x}}\in\mathscr{X}} \frac{f_i(\tilde{\boldsymbol{x}}) - f_i(\boldsymbol{x}) - \langle \nabla f_i(\boldsymbol{x}),\tilde{\boldsymbol{x}}-\boldsymbol{x}\rangle}{\frac{1}{2}\|\tilde{\boldsymbol{x}}-\boldsymbol{x}\|_2^2} \\
&= m.
\end{aligned}$$

Since computing the minimum here is difficult and is generally a non-convex problem, we can instead look at an approximate method. Suppose that we have N points
$\boldsymbol{x}(1),\ldots,\boldsymbol{x}(N)$. Then for any two distinct points $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$, we have

$$m \leq \frac{f_i(\boldsymbol{x}(i)) - f_i(\boldsymbol{x}(j)) - \langle \nabla f_i(\boldsymbol{x}(j)),\boldsymbol{x}(i)-\boldsymbol{x}(j)\rangle}{\frac{1}{2}\|\boldsymbol{x}(i)-\boldsymbol{x}(j)\|_2^2}.$$

This suggests the estimate

$$\hat{m}_i \triangleq \min_{i\neq j} r(\boldsymbol{x}(i),\boldsymbol{x}(j)) \tag{54}$$

for the strong convexity parameter. Then we have

$$\begin{aligned}
&\mathbb{E}[\hat{m}_i \mid \mathscr{F}_{i-1}] \\
&= \mathbb{E}\left[\min_{i\neq j} r(\boldsymbol{x}(i),\boldsymbol{x}(j)) \,\middle|\, \mathscr{F}_{i-1}\right] \\
&\leq \min_{i\neq j} \mathbb{E}\left[r(\boldsymbol{x}(i),\boldsymbol{x}(j)) \,\middle|\, \mathscr{F}_{i-1}\right] \\
&\leq \min_{i\neq j} \frac{f_i(\boldsymbol{x}(i)) - f_i(\boldsymbol{x}(j)) - \langle \nabla f_i(\boldsymbol{x}(j)),\boldsymbol{x}(i)-\boldsymbol{x}(j)\rangle}{\frac{1}{2}\|\boldsymbol{x}(i)-\boldsymbol{x}(j)\|_2^2}.
\end{aligned}$$

It is difficult to compare this estimate to m exactly. All we can say is that

$$m \leq \min_{i\neq j} \frac{f_i(\boldsymbol{x}(i)) - f_i(\boldsymbol{x}(j)) - \langle \nabla f_i(\boldsymbol{x}(j)),\boldsymbol{x}(i)-\boldsymbol{x}(j)\rangle}{\frac{1}{2}\|\boldsymbol{x}(i)-\boldsymbol{x}(j)\|_2^2}$$

as well.

13.2 Estimating Gradient Parameters

We seek (A,B) such that

$$\mathbb{E}\|\boldsymbol{g}(\boldsymbol{x},\boldsymbol{z})\|_2^2 \leq A + B\|\boldsymbol{x}-\boldsymbol{x}^*\|_2^2.$$

Suppose that our functions have Lipschitz continuous gradients with modulus M, and we construct estimates of the modulus $\hat{M}_i$ analogous to eq. (53) or eq. (54) by replacing the min with a max. Suppose that we select N points $\boldsymbol{x}(1),\ldots,\boldsymbol{x}(N) \in \mathscr{X}$. We want to find A and B such that

$$\mathbb{E}\|\boldsymbol{g}(\boldsymbol{x}(j),\boldsymbol{z})\|_2^2 \leq A + B\|\boldsymbol{x}(j)-\boldsymbol{x}^*\|_2^2.$$

By the Lipschitz gradient assumption, we have

$$\|\boldsymbol{x}-\boldsymbol{x}^*\|_2 \geq \frac{1}{M}\|\nabla f(\boldsymbol{x})\|_2.$$

Therefore, the following implication holds

$$\mathbb{E}\|\boldsymbol{g}(\boldsymbol{x}(j),\boldsymbol{z})\|_2^2 \leq A + \frac{B}{M^2}\|\nabla f(\boldsymbol{x}(j))\|_2^2 \;\Rightarrow\; \mathbb{E}\|\boldsymbol{g}(\boldsymbol{x}(j),\boldsymbol{z})\|_2^2 \leq A + B\|\boldsymbol{x}(j)-\boldsymbol{x}^*\|_2^2.$$

We look for (A,B) such that

$$\mathbb{E}\|\boldsymbol{g}(\boldsymbol{x}(j),\boldsymbol{z})\|_2^2 \leq A + \frac{B}{M^2}\|\nabla f(\boldsymbol{x}())\|_2^2 \quad j=1,\ldots,N.$$

Define

$$s_i(j) \triangleq \frac{1}{K_i}\sum_{k=1}^{K_i}\|\boldsymbol{g}_i(\boldsymbol{x}(j),\boldsymbol{z}_i(k))\|_2^2$$

and

$$d_i(j) \triangleq g_i(j) - \frac{1}{K_i-1}\sum_{k=1}^{K_i}\left\|\boldsymbol{g}_i(\boldsymbol{x}(j),\boldsymbol{z}_i(k)) - \frac{1}{K_i}\sum_{p=1}^{K_i}\boldsymbol{g}_i(\boldsymbol{x}(p),\boldsymbol{z}_i(p))\right\|_2^2.$$

We want to find (A,B) such that

$$s_i(j) \leq A + \frac{B}{(\hat{M}_{i-1}+t_{i-1})^2}d_i(j) \quad j=1,\ldots,N.$$

Suppose that we are given a function $\phi(A,B)$ that controls the size of (A,B). For example, we may have $\phi(A,B) = \frac{1}{2}A^2 + \frac{1}{2}B^2$ or $\phi(A,B) = \lambda A^2 + (1-\lambda)B^2$ with $0 \leq \lambda \leq 1$. We solve

$$\begin{aligned} \underset{\tilde{A}_i,\tilde{B}_i}{\text{minimize}} \quad & \phi(\tilde{A}_i,\tilde{B}_i) \\ \text{subject to} \quad & s_i(j) \leq \tilde{A}_i + \frac{\tilde{B}_i}{(\hat{M}_{i-1}+t_{i-1})^2}d_i(j), \quad j=1,\ldots,N \\ & \tilde{A}_i \geq 0\ ,\ \tilde{B}_i \geq 0 \end{aligned} \tag{55}$$

to generate approximate $(\tilde{A}_i,\tilde{B}_i)$.

13.3 Combining One Step Estimates

One issue in parameter estimation is that there may be some dependencies among the various estimates which need to be accounted for. For example, the estimates for (A,B) in eq. (55) depend on estimates for the Lipschitz modulus M. We show that this does not impact our estimation process using Lemma 13 and Lemma 14. First, we present a result showing that if we plug in the true parameters that our estimates work.

Lemma 13. *Suppose that we estimate ϕ^* by averaging the estimates $\phi_i(\pi^*)$ where π^* are the true parameters on which the estimate ϕ_i depends and the following conditions hold:*

1. $|\phi_i(\pi^*)| \leq C$

2. $\mathbb{E}[\phi_i(\pi^*) \mid \mathscr{F}_{i-1}] \geq \phi^*$

3. $\sum_{n=1}^{\infty} \exp\left\{-\frac{2nt_n^2}{C^2}\right\} < +\infty$

Then for all n large enough, it holds that

$$\frac{1}{n}\sum_{i=1}^{n} \phi_i(\pi^*) + t_n \geq \phi^*$$

almost surely.

Proof. Since $|\phi_i(\pi^*)| \leq C$, by applying Lemma 4 it holds that

$$\mathbb{E}\left[e^{s(\phi_i(\pi^*) - \mathbb{E}[\phi_i(\pi^*) \mid \mathscr{F}_{i-1}])} \mid \mathscr{F}_{i-1}\right] \leq \exp\left\{\frac{1}{2}\frac{C^2}{4}s^2\right\}.$$

Then by Lemma 3, it holds that

$$\begin{aligned}\mathbb{P}\left\{\frac{1}{n}\sum_{i=1}^{n}\phi_i(\pi^*) < \phi^* - t_n\right\} &= \mathbb{P}\left\{\frac{1}{n}\sum_{i=1}^{n}\phi_i(\pi^*) < \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}\left[\phi_i(\pi^*) \mid \mathscr{F}_{i-1}\right] - t_n\right\} \\ &\leq \exp\left\{-\frac{2nt_n^2}{C^2}\right\}.\end{aligned}$$

Since it holds that

$$\sum_{n=1}^{\infty}\mathbb{P}\left\{\frac{1}{n}\sum_{i=1}^{n}\phi_i(\pi^*) < \phi^* - t_n\right\} \leq \sum_{n=1}^{\infty}\exp\left\{-\frac{2nt_n^2}{C^2}\right\} < +\infty$$

by the Borel-Cantelli lemma, it follows that for all n large enough

$$\frac{1}{n}\sum_{i=1}^{n}\phi_i(\pi^*) + t_n \geq \phi^*.$$

□

Since our estimates of m do not depend on any parameters π, this lemma shows that the averaged estimate eventually lower bounds m. Similar reasoning holds for the Lipschitz modulus M. 13 also shows that estimates of (A, B) upper bound the true quantities provided that the true value of m and M are plugged in. We bootstrap from this result to show that the estimates of A and B upper bound the exact quantities using Lemma 14. Before proceeding, note that random variables X_n are $o_{\mathbb{P}}(1)$ if

$$\lim_{n\to\infty}\mathbb{P}\{|X_n| \geq t\} = 0 \quad \forall t > 0.$$

Lemma 14. *Suppose that we estimate ϕ^* by averaging the estimates $\phi_i(\pi_i)$ where π_i are the estimates of the parameters on which the estimate ϕ_i depends and the following hold:*

1. $|\phi_i(\pi)| \leq C$
2. *For all n large enough $\pi_n \geq \pi^*$ almost surely*
3. $\pi \leq \tilde{\pi} \;\Rightarrow\; \phi_i(\pi) \leq \phi_i(\tilde{\pi})$
4. *For appropriate sequences t_n, $\frac{1}{n}\sum_{i=1}^{n}\phi_i(\pi^*) + t_n \geq \phi^*$*

Then for all n large enough, it holds that

$$\frac{1}{n}\sum_{i=1}^{n}\phi_i(\pi_i) + t_n \geq \phi^* + o_{\mathbb{P}}(1)$$

almost surely.

Proof. There exists a finite almost surely random variable $\tilde{N}$ such that

$$n \geq \tilde{N} \;\; \Rightarrow \;\; \pi_i \geq \pi^*.$$

It holds that

$$\begin{aligned}
\frac{1}{n}\sum_{i=1}^{n}\phi_i(\pi_i) &= \frac{1}{n}\sum_{i=1}^{\tilde{N}-1}\phi_i(\pi_i) + \frac{1}{n}\sum_{i=\tilde{N}}^{n}\phi_i(\pi_i) \\
&\geq \frac{1}{n}\sum_{i=1}^{\tilde{N}-1}\phi_i(\pi_i) + \frac{1}{n}\sum_{i=\tilde{N}}^{n}\phi_i(\pi^*) \\
&= \frac{1}{n}\sum_{i=1}^{\tilde{N}-1}(\phi_i(\pi_i) - \phi_i(\pi^*)) + \frac{1}{n}\sum_{i=1}^{n}\phi_i(\pi^*).
\end{aligned}$$

By the boundedness of $\phi(\pi)$, this implies that

$$\begin{aligned}
\frac{1}{n}\sum_{i=1}^{n}\phi_i(\pi_i) + t_n &= \left(\frac{1}{n}\sum_{i=1}^{n}\phi_i(\pi^*) + t_n\right) + o_{\mathbb{P}}(1) \\
&\geq \phi^* + o_{\mathbb{P}}(1).
\end{aligned}$$

□

This result proves that estimates for (A, B) work as a result of the estimates for m and M working.

14 Experiment

We look at an example of the IHP tracking criterion using the same example as table 3. We also compare against the Kalman filter when the minimizers change according to a random walk model.

14.1 IHP Tracking Criterion

Figure 6 plots r vs ε for several values of K by applying the IHP algorithm. The IHP bounds appear to be loose in general as we need fairly large values of K to get non-trivial bounds for reasonable ε and small r. The looseness of these bounds is not surprising, since we are only using the first moment of the tracking error to bound.

We choose K_n by targeting $t = 0.1$ and $r = 0.25$. Figure 7 shows the empirical probability that $f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t$. As mentioned above, we can compute $f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*)$ exactly, so we can calculate the fraction of the time that the loss violates the $t = 0.1$ constraint. The empirical probability that $f_n(\boldsymbol{x}_n) - f_n(\boldsymbol{x}_n^*) > t$ satisfies our target value of $r = 0.25$ to within the error bars.

14.2 Kalman Filter Comparison

We now consider a slight modification of our model, so that we can apply the Kalman filter. As mentioned above, since we assume that $\boldsymbol{\eta}_n$ is generated as a deterministic sequence, we cannot apply the Kalman filter. In this section, we instead assume that

$$\boldsymbol{\eta}_n - \boldsymbol{\eta}_{n-1} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \boldsymbol{I})$$

and $\boldsymbol{\eta}_1$ is fixed. Then it holds that

$$\|\boldsymbol{\eta}_n - \boldsymbol{\eta}_{n-1}\|_{L_2} \leq \sigma \triangleq \rho.$$

We satisfy (2) and use the estimate of ρ in (24).

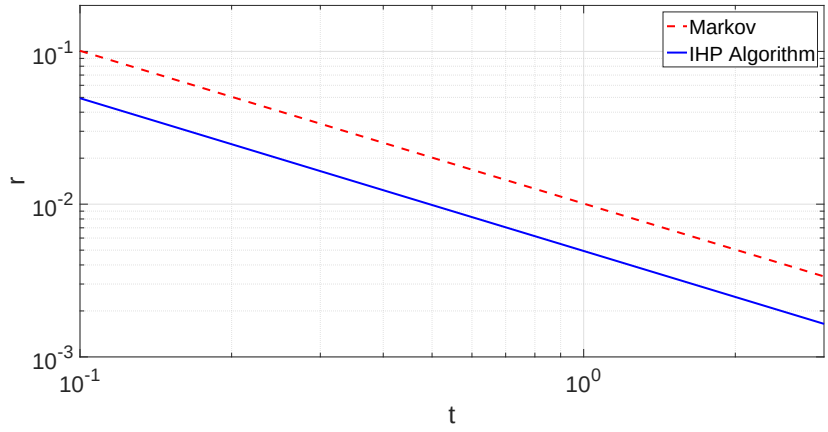

Figure 6: In high probability criterion: r vs t

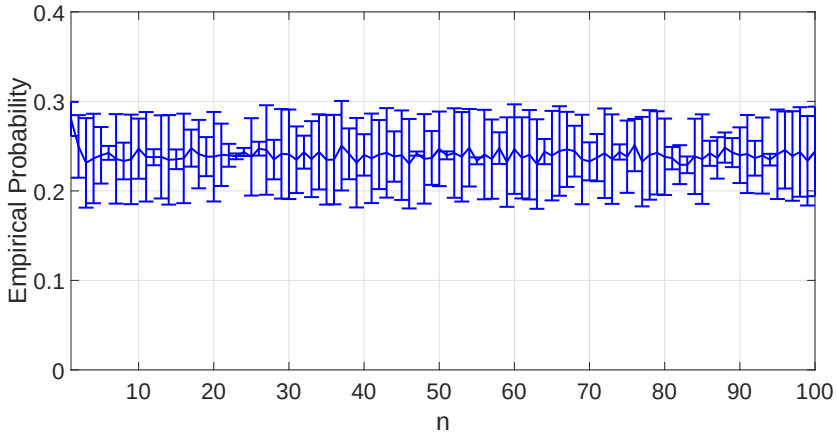

Figure 7: Empirical Probability

To apply the Kalman filter, we take $\boldsymbol{\eta}_n$ to be the state of the system. The state evolution equation is given by

$$\boldsymbol{\eta}_n(k) = \begin{cases} \boldsymbol{\eta}_1(1), & \text{fixed} \\ \boldsymbol{\eta}_{n-1}(K_{n-1}) + \zeta_n, & k = 1 \\ \boldsymbol{\eta}_n(k-1), & 1 < k \leq K_n \end{cases}$$

with $\zeta_n \sim \mathcal{N}(\mathbf{0}, \sigma^2 \boldsymbol{I})$. The observation equation is given by the pair $(\boldsymbol{w}_n(k), y_n(k))$ with

$$y_n(k) = \boldsymbol{\eta}_n^\top(k)\boldsymbol{w}_n(k) + e_n(k).$$

Let $\hat{\boldsymbol{\eta}}_n(k|\tilde{k})$ be the estimate of $\boldsymbol{\eta}_n$ at time k of epoch n given all the information up to $\tilde{k}$ with $k \geq \tilde{k}$. Let $P_n(k|\tilde{k})$ be the estimate of the covariance. The prediction equations for the state estimate and covariance estimate are given by [36]

$$\begin{aligned} \hat{\boldsymbol{\eta}}_n(k|k-1) &= \hat{\boldsymbol{\eta}}_n(k-1|k-1) \\ P_n(k|k-1) &= P_n(k-1|k-1) + \sigma^2 \boldsymbol{I} \mathbb{1}_{\{k=1\}}. \end{aligned} \tag{56}$$

The update equations are given by

$$\begin{aligned}
\hat{\boldsymbol{\eta}}_n(k|k) &= \hat{\boldsymbol{\eta}}_n(k|k-1) + \boldsymbol{G}_n(k)(y_n(k) - \hat{\boldsymbol{\eta}}_n^\top(k|k-1)\boldsymbol{w}_n(k)) \\
P_n(k|k) &= (\boldsymbol{I} - G_n(k)\boldsymbol{w}_n^\top(k))P_n(k|k-1) \\
G_n(k) &= P_n(k|k-1)\boldsymbol{w}_n(k)\left(\sigma_e^2 + \boldsymbol{w}_n^\top(k)P_n(k|k-1)\boldsymbol{w}_n(k)\right)^{-1}
\end{aligned}$$

where $G_n(k)$ is the Kalman gain. We have the initial conditions

$$\begin{aligned}
\hat{\boldsymbol{\eta}}_n(1|0) &= \hat{\boldsymbol{\eta}}_{n-1}(K_{n-1}|K_{n-1}) \\
P_n(1|0) &= P_{n-1}(K_{n-1}|K_{n-1}).
\end{aligned}$$

Figure 8 shows a comparison of the Kalman filter against our SGD based approach both with exact and mismatched parameters for the Kalman filter. Table 3 uses the technique from (33) to estimate the mean criterion for all three methods. The Kalman filter receives the number of samples K_n chosen by the SGD approach. With correct parameters for the Kalman filter, both methods achieve similar performance, but the SGD method is able to control its desired accuracy. With incorrect parameters, the Kalman filter's performance is considerably worse.

Method	Mean Criterion Estimate
Direct Estimate	$1.9 \times 10^{-2} \pm 1.1 \times 10^{-3}$
Kalman Filter	$1.8 \times 10^{-2} \pm 5.7 \times 10^{-3}$
Kalman Filter - Mismatch	$9.4 \times 10^{-2} \pm 5.3 \times 10^{-2}$

Table 3: Kalman Filter Comparison

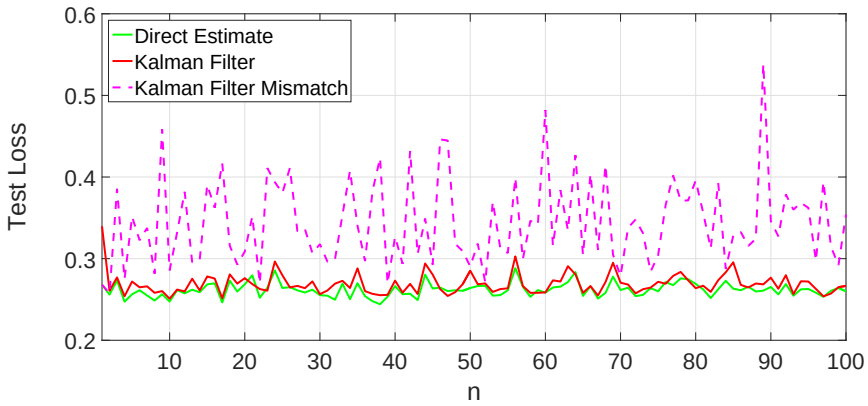

Figure 8: Comparison of Kalman filter and our approach